



RED HAT TECH UPDATE APRIL 2018

Peter Gustafsson
Solutions Architect
Red Hat Sweden

Johan Odell
Solutions Architect
Red Hat Sweden

Johan Robinson
Specialist Solutions Architect - Storage
Red Hat EMEA

AGENDA

Red Hat Tech Update

- Red Hat Satellite 6.3 New Features
- Gluster Storage
- Ceph Storage
- RHHI
- ~~Red Hat CloudForms 4.6~~

The DEMO-environment is
proudly brought to you by



Red Hat Satellite 6.3 New Features

Key Focus Areas



Content management




Supportability



System provisioning



Security & user access



Configuration management



Usability

PACKAGING

SIMPLIFIED PRICING AND PACKAGING

Get the most out of your Red Hat infrastructure investment

OLD MODEL

RED HAT[™] SATELLITE

Purchase each of the following:

- Red Hat Satellite Server
- Red Hat Satellite Capsule Server
- Smart Management Add-On

RED HAT[™] INSIGHTS

Purchase the following separately:

- Red Hat Insights Add-On

NEW MODEL


RED HAT[™]
SATELLITE + RED HAT[™]
INSIGHTS

Receive all Satellite and Insights
elements (servers & clients) with the
purchase of Smart Management

CONTENT MANAGEMENT

IMPROVED CONTENT DOWNLOAD POLICIES AND SYNCHRONIZATION

Better control over download policies using Lazy Sync



Lazy Sync Modes:

All modes applicable to the Red Hat Satellite Server and Red Hat Satellite Capsule Servers. Capsules can now set a content download policy independent of the Satellite settings.

On demand (default)

Retrieves content only when systems request it.

Background

Allows repositories to be published and made available to systems while the content downloads asynchronously.

Immediate

Downloads all content before publishing the repository.

CUSTOM FILE-TYPE REPOSITORIES

Unify the delivery of managed content



NEW FEATURE

Create repositories containing custom file types with new repository management tooling and import those repositories into Red Hat Satellite.

Version management

Custom file types can be added to content views and versioned similar to other content types (RPMs, ISOs, etc).

SYSTEM PROVISIONING

PULL TEMPLATES FROM GIT

Improved ability to manage provisioning templates, similar to content versioning



Manage and Deploy

Edit templates and clone them using GIT branching and versioning, or any version control system.

CONFIGURATION MANAGEMENT

SATELLITE AND ANSIBLE TOWER INTEGRATION

Documented best practices to help optimize use of both products



By integrating Red Hat Satellite with Red Hat Ansible[®] Tower, administrators can now perform the following functions:

Dynamic inventory

Allows Ansible Tower to use Satellite as a dynamic inventory source.

Provisioning callbacks

Allows systems provisioned via Satellite to “callback” to Ansible Tower so that playbook runs can happen post-provisioning.

SUPPORTABILITY

PLATFORM SUPPORTABILITY ENHANCEMENTS



NEW FEATURES



puppet



Amazon EC2

Puppet 3.8 & 4

Customers have the freedom to use Puppet 4 within Red Hat Satellite in lieu of (or in conjunction with) their existing Puppet 3.8 deployment.

Feature Overview: <https://access.redhat.com/articles/3358711>

Unified Extensible Firmware Interface support (UEFI)

Provides mechanisms to allow for support of new technologies, improved development, and enhanced customer experience during the time before the operating system loads.

Amazon Web Services (AWS)

AWS EC2 is now a supported platform for Red Hat Satellite & Red Hat Satellite Capsule Server).

Feature Overview: <https://access.redhat.com/articles/3358631>

SATELLITE AND INSIGHTS INTEGRATION

Deeper integration with Red Hat Insights



By integrating Red Hat Satellite with Red Hat Insights, administrators can now perform the following functions:

Dashboard Widgets

New widgets for Insights Actions and Insights Risk Summary


Insights Planner

Build out an Insights Planner Playbook for use with Ansible or Ansible Tower from Red Hat Satellite

SECURITY & USER ACCESS

NEW “ORG ADMIN” ROLE

Effectively limit administrator scope within Red Hat Satellite management



NEW FEATURE

The org admin role is intended to scope an administrator such that their reach within the full extent of the Satellite management plane is constrained:

Multi-tenant


No visibility into other organizations. No exposure to the existence of other organizations.

Org admin role

Cannot access existing Satellite users who already assigned other alternative organizations.

OPENS CAP TAILORING FILES

Freedom to customize OpenSCAP policy application across systems



**NEW
FEATURE**

Tailoring files

Change the behavior of an existing OpenSCAP policy without having to fork or rewrite the entire policy.

USABILITY

NOTIFICATION DRAWER

Quick access to occurring events across Red Hat Satellite



NEW FEATURE

Quick Access

Provides fast access to important events

Easy interaction


Gives admins the opportunity to clear or interact with the information provided in the event.

New Functionality

Limited events use the Notification Drawer today. Additional events will move to the notification drawer in the future

FUTURE-DATED SUBSCRIPTIONS

Avoid “red” status by applying subscriptions in advance



NEW FEATURE

Customers can now import purchased Red Hat subscriptions **before the actual start date** – and allocate them as needed.

Continual coverage


Systems can be configured to have overlap between expiring and new subscriptions, ensuring continual access to content.

Future-dated subscriptions

Do not provide access to content until their start date.

SATELLITE CLONE

Ability to clone existing Satellite Server to a new host
Provides flexibility with Red Hat product upgrades and migrations

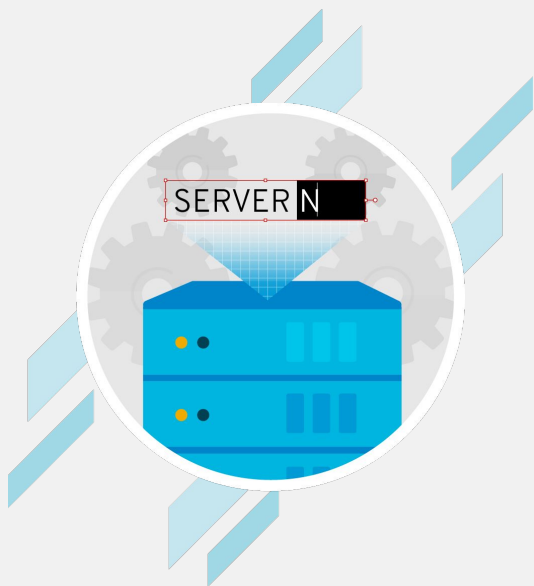


Easy Cloning

Provides a means for the customer to clone their Red Hat Satellite to assist with server migration, or to upgrade from one version of Red Hat Enterprise Linux to another.

SATELLITE RENAMING TOOL

Change the Satellite hostname while changing configuration



Automated reconfiguration of the server

Change the hostname of Red Hat Satellite or Red Hat Satellite Capsule Server while updating the various configurations needed to accomplish the change.

VIRT-WHO CONFIGURATION WIZARD

Simplified configuration and deployment of virt-who



Easy virt-who configuration

The virt-who wizard aids in running and configuring the virt-who daemon on any Red Hat Enterprise Linux Server being managed by Red Hat Satellite.

NEXT STEPS & RESOURCES

NEXT STEPS

Review all Satellite 6.3 feature overviews and videos:

<https://access.redhat.com/blogs/1169563>

Use the Upgrade Helper to plan your upgrade-

<https://access.redhat.com/labs/satelliteupgradehelper/>

Attention Satellite 5 customers:

- End of support for versions 5.7 & below is January 31, 2019 - access to content will not be provided after this date!
 - **Recommended action ASAP: Upgrade to Satellite 5.8** (supported thru May 2020)
 - Consider moving to Satellite 6, when ready

RESOURCES

Red Hat Satellite Product page - <http://redhat.com/satellite>

Red Hat Satellite Customer Portal -
<https://access.redhat.com/products/red-hat-satellite>

Red Hat Satellite Documentation -
https://access.redhat.com/documentation/en-us/red_hat_satellite/

Red Hat Gluster Storage

Red Hat Gluster Storage

RED HAT® GLUSTER STORAGE

TARGET USE CASES

Container Storage

- Persistent storage
- Registry storage

Enterprise File Sharing

- Media streaming
- Active Archives

Enterprise Virtualization

Flexible file storage for petabyte-scale workloads

- Purpose-built as a scale-out file store with a straightforward architecture suitable for public, private, and hybrid cloud
- Simple to install and configure, with a minimal hardware footprint
- Offers mature NFS, SMB and object (Swift) and native glusterfs-fuse interfaces



Media



Machine Data



CCTV



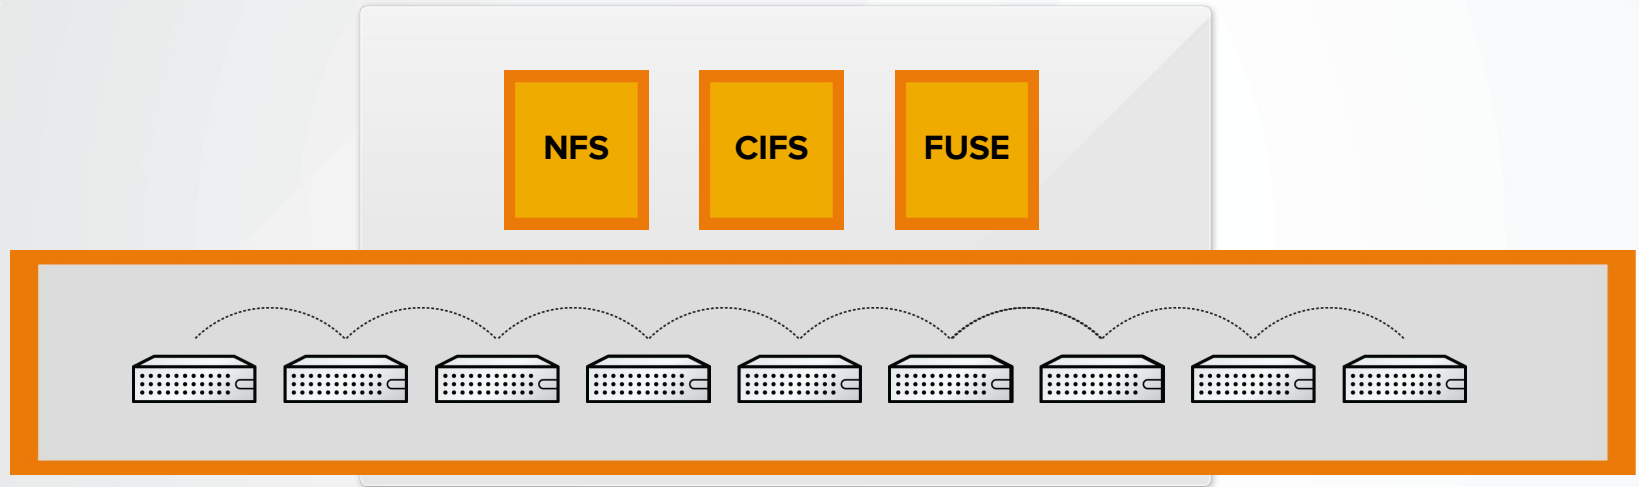
Archive



Documents


GLUSTER ARCHITECTURE

- Distributed scale out share nothing storage using industry standard hardware



Aggregates systems to one cohesive unit
and presents using common protocols

System anatomy



HW

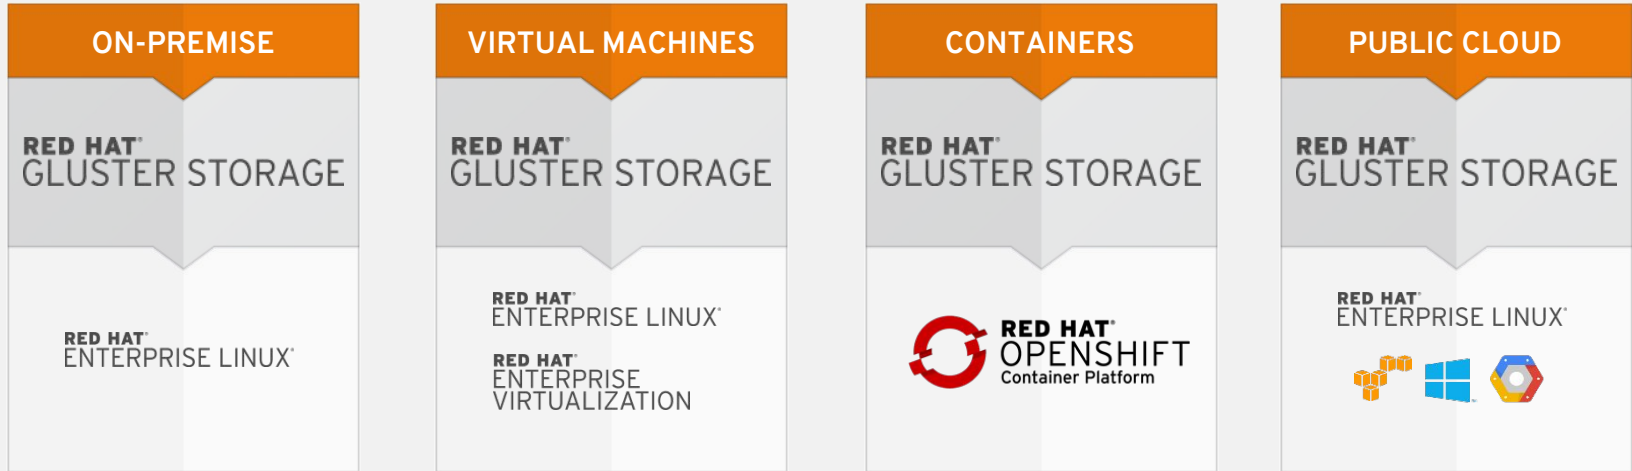
- x86 standard server RHEL certified
- RAID supported RAID6, RAID10, RAID5
- Full Flash SSD and HDD architectures
- Ethernet 10Gb/s, 40Gb/s, Infiniband supported

SW

- Self-healing and data integrity check
- HA with mirroring and data replication
- Native Scale-Out and Scale-Up
- Geo and In-system replication

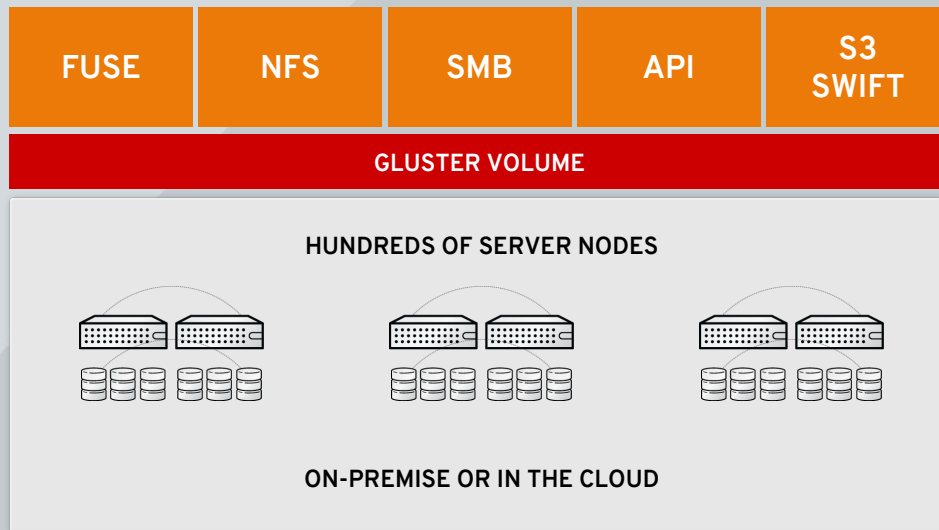
Flexible Deployment with RHGS

- Same software bits across on-prem, VMs, containers and all three public clouds
- Applications can be ported across deployments without expensive re-writes
- Close integration with RHEL, RHV, and OpenShift



MULTI-PROTOCOL ACCESS

Primarily accessed as scale-out file storage with optional APIs, Swift or S3 object




Red Hat Gluster Storage

What's new...

ARBITER VOLUMES

Cost effective alternative to 3-way replication

- **REDUCTION IN FOOTPRINT & TCO**
Huge capacity savings
- **COST-EFFECTIVE DATA INTEGRITY**
Integrity of 3-way w/o 3x capacity
- **MULTIPLE DEPLOYMENT MODELS**
Dedicated or Daisy chained
- **USE CASES**
Backup/archiving, HCI, Stretched clusters



• ERASURE CODING VOLUME ENHANCEMENTS

The following new variants are now supported:

- 10 bricks with redundancy level 2 (8 + 2)
- 20 bricks with redundancy level 4 (16 + 4)

(in addition to the already supported 4+2, 8+3 and 8+4 configurations)

- **PERFORMANCE IMPROVEMENTS**

“*find*” command on a volume improved when BitRot detection is used

Parallel readdirp support

readdirp fops are sent parallelly to all the bricks. This enhances the performance for find and a recursive listing of small directories.

Enhancement to glusterfind command

provides a query sub command that provides a list of changed files

- **NFS GANESHA**

Dynamic update of export configuration options

Most NFS-Ganesha export configuration options can be updated dynamically during normal operation without needing to export and re-export the volume.

- **GEO REPLICATION STATUS CMD**

The detailed geo-replication status command no longer requires master volume, slave host, and slave volume.

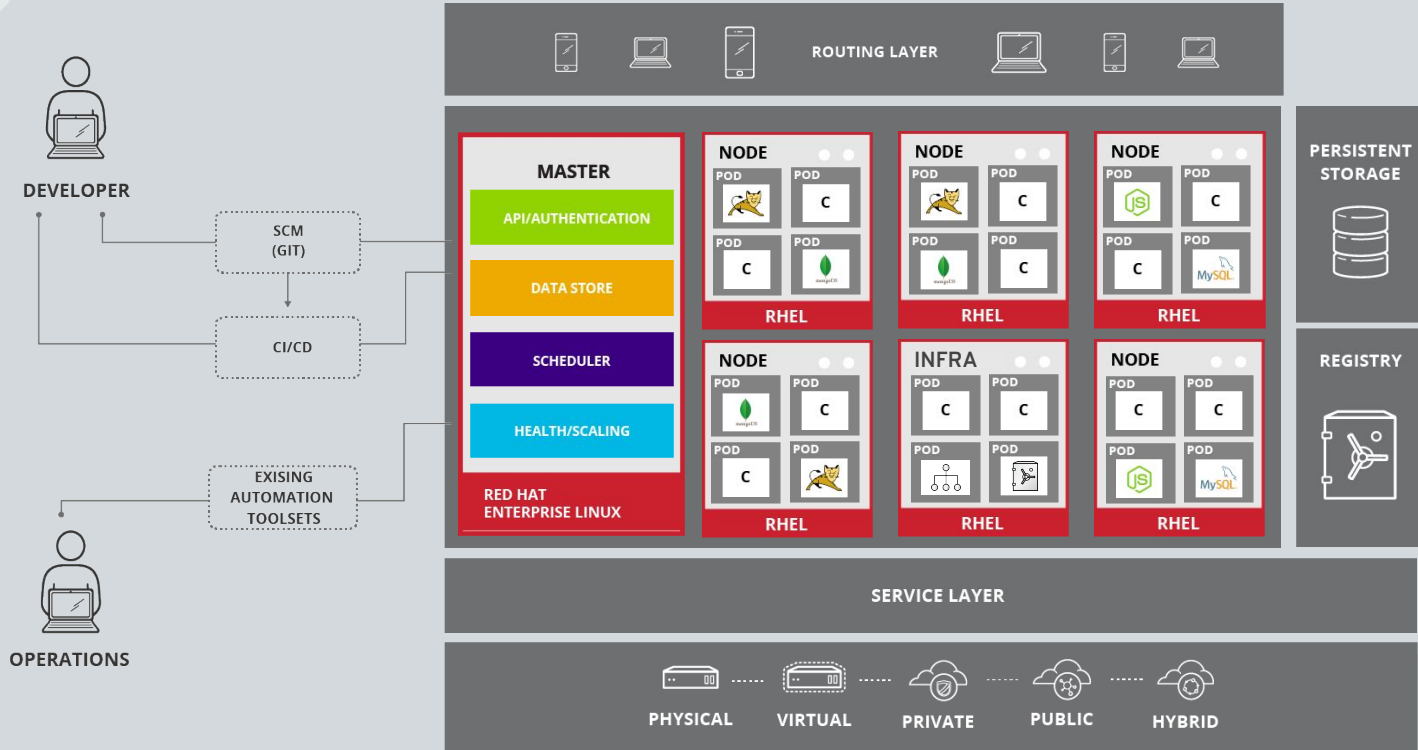
It can be executed with or without these additional details.

gluster volume geo-replication status detail

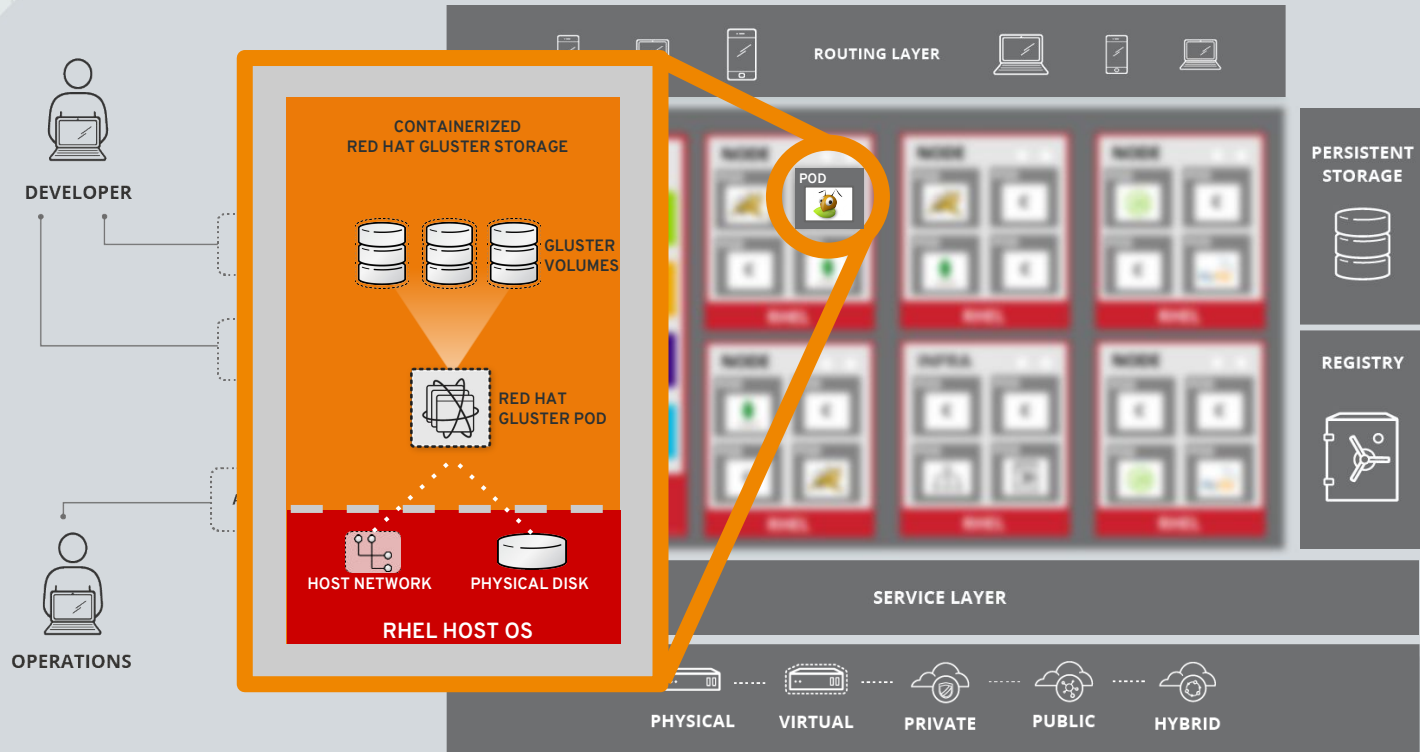
- **ENHANCED REBALANCE STATUS**

The command *gluster volume volname rebalance status* now provides an estimate of the time left to rebalance completion. Calculations are based on each brick having its own file system partition.

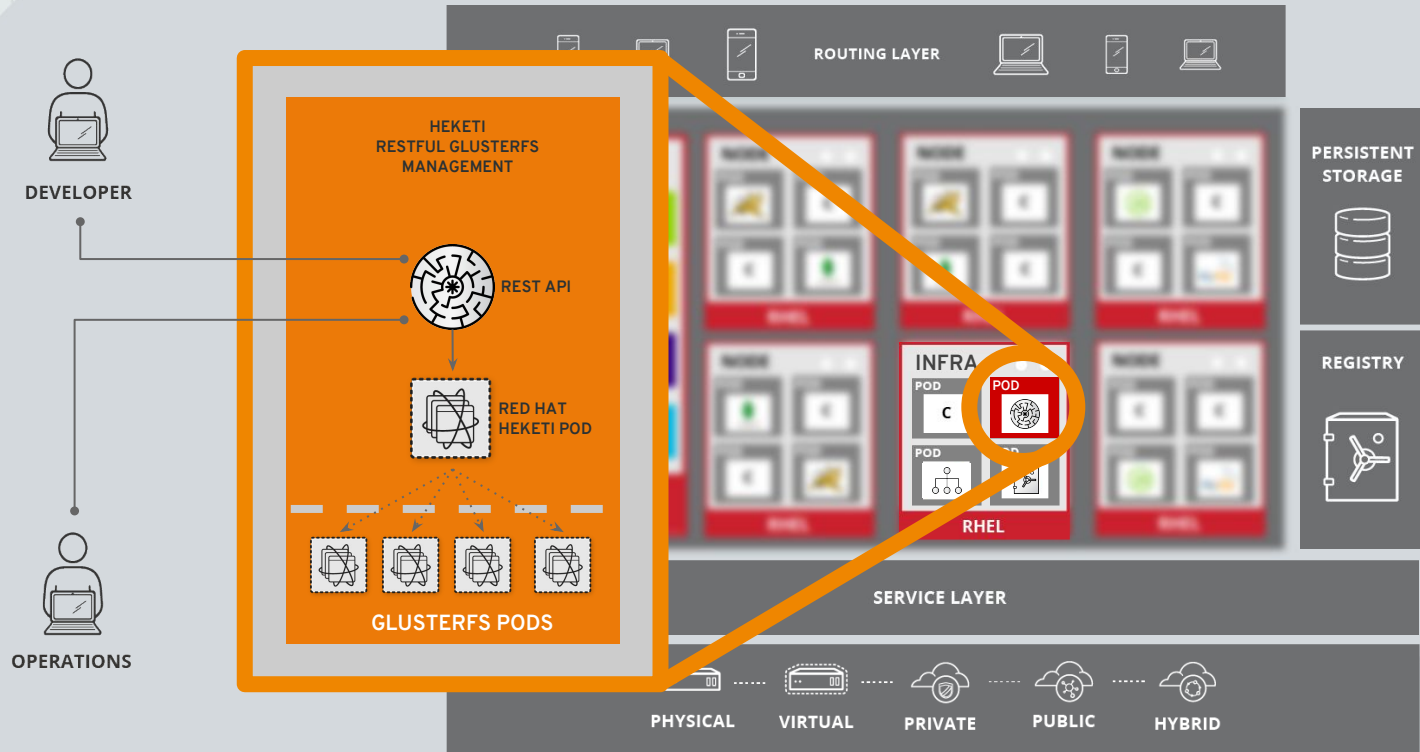
CONTAINER-NATIVE STORAGE ON OPENSHIFT



CONTAINER-NATIVE STORAGE ON OPENSHIFT



CONTAINER-NATIVE STORAGE ON OPENSHIFT



CONTAINER-NATIVE STORAGE ON OPENSHIFT

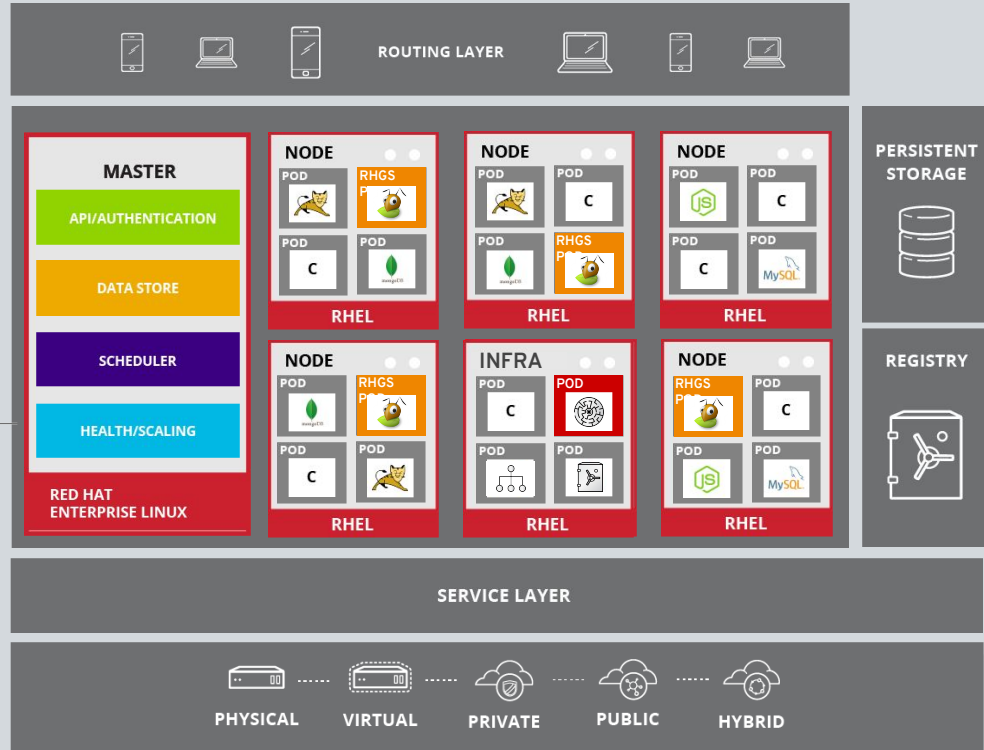


DEVELOPER

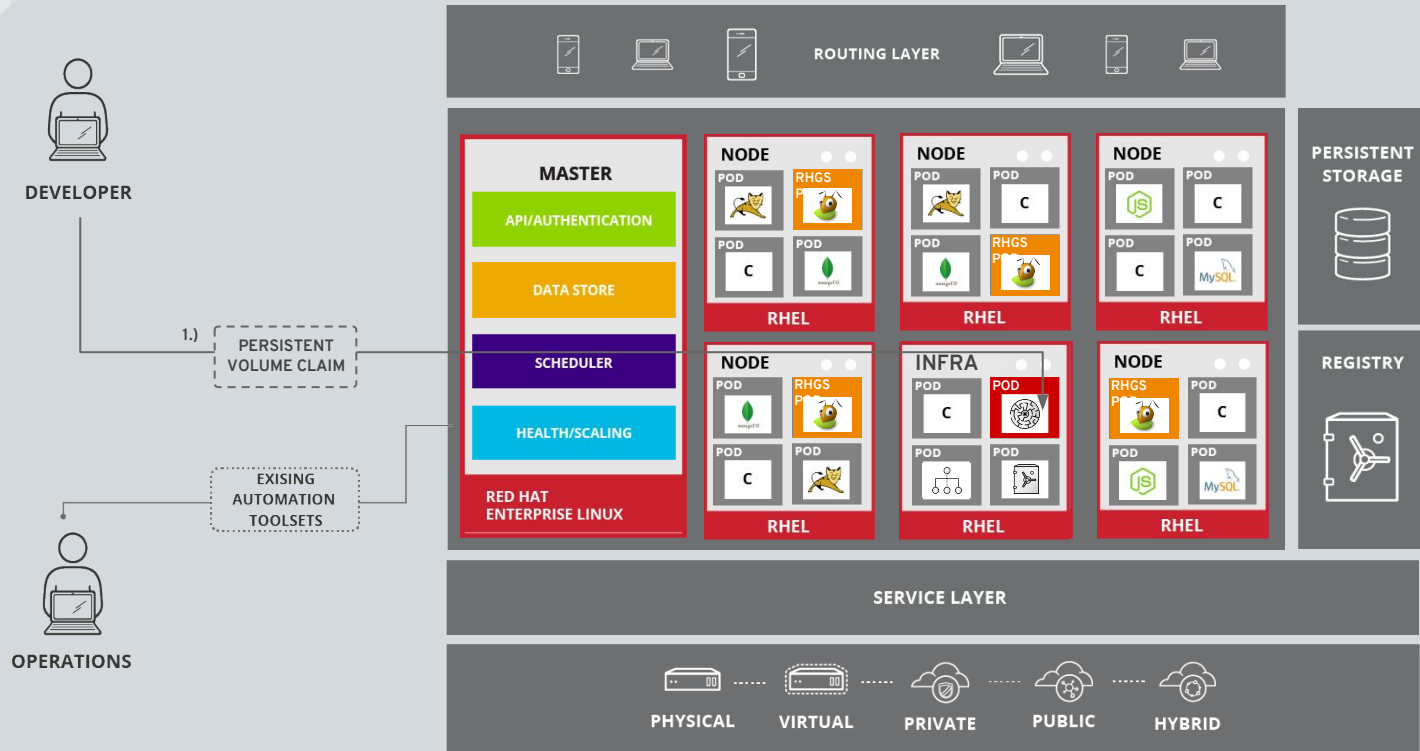


OPERATIONS

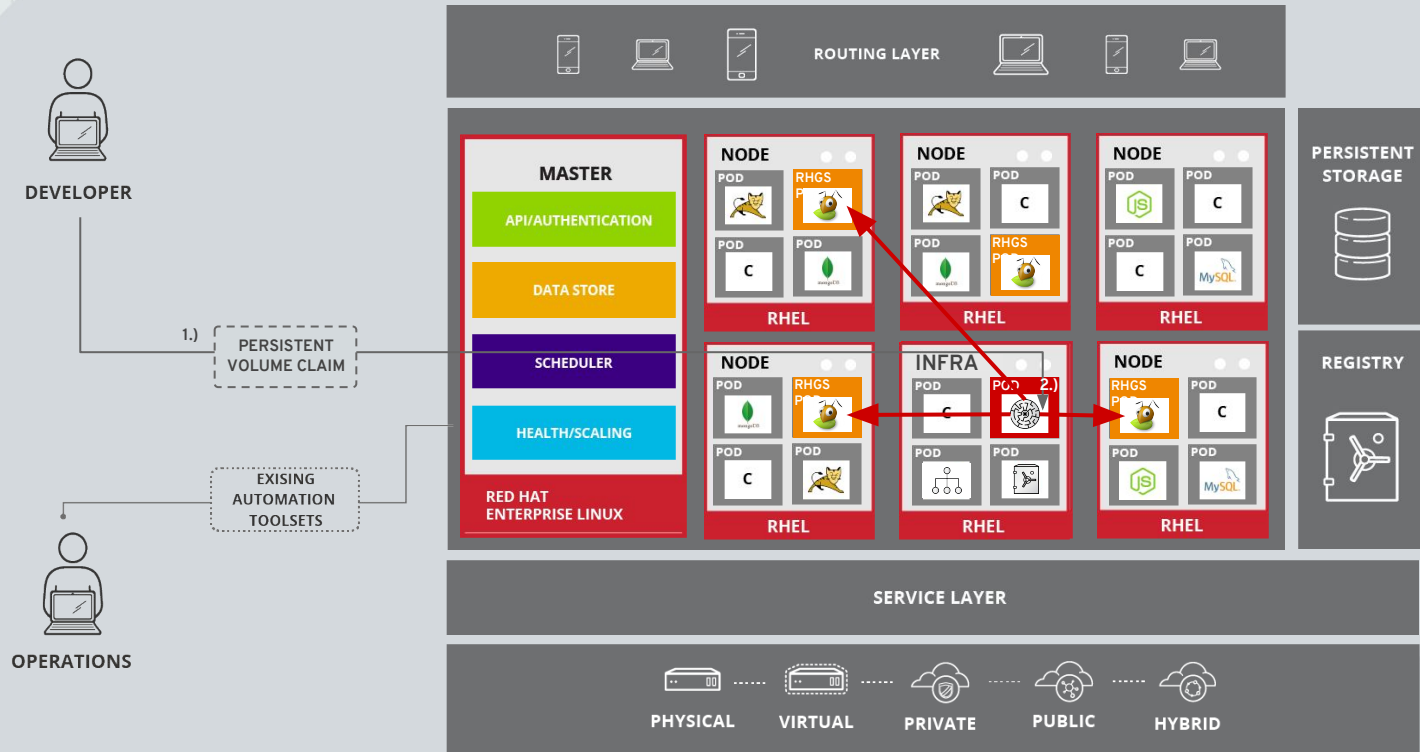
EXISTING
AUTOMATION
TOOLSETS



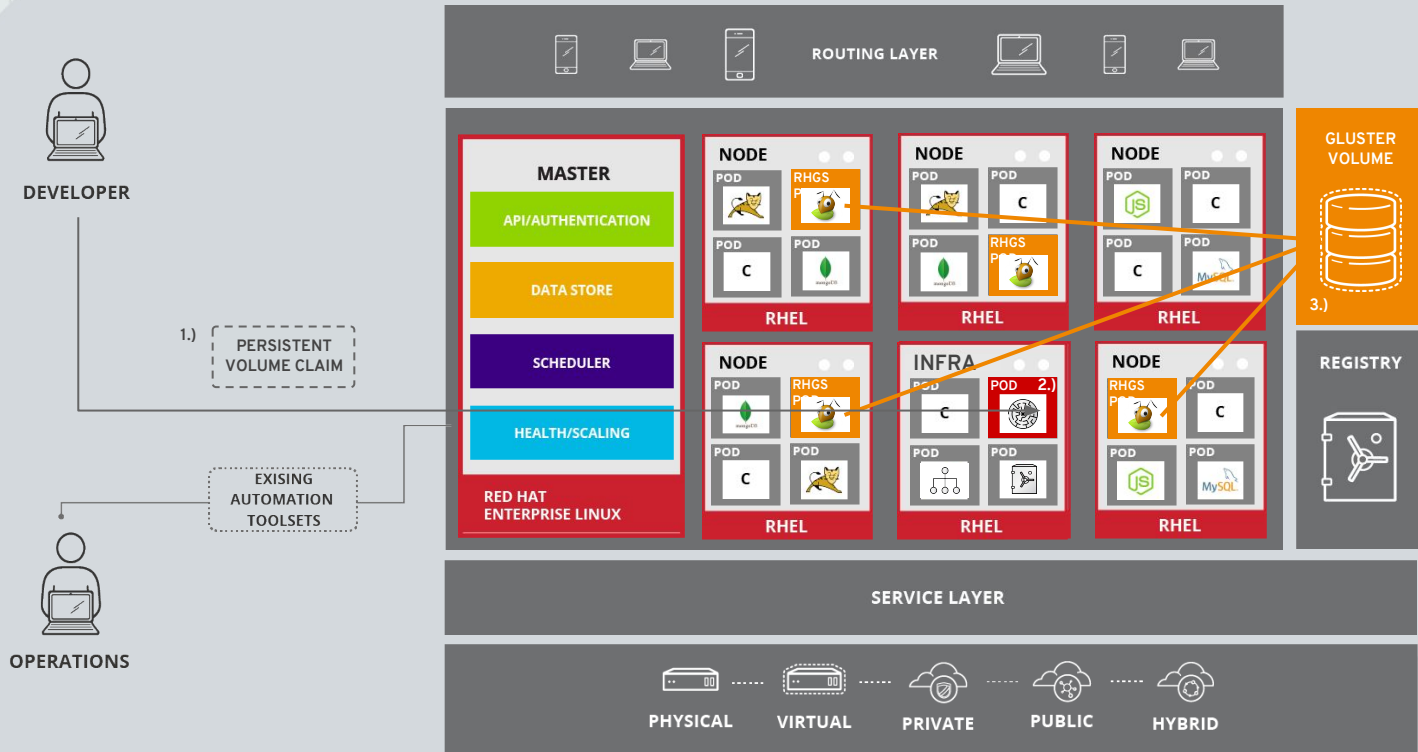
CONTAINER-NATIVE STORAGE ON OPENSHIFT



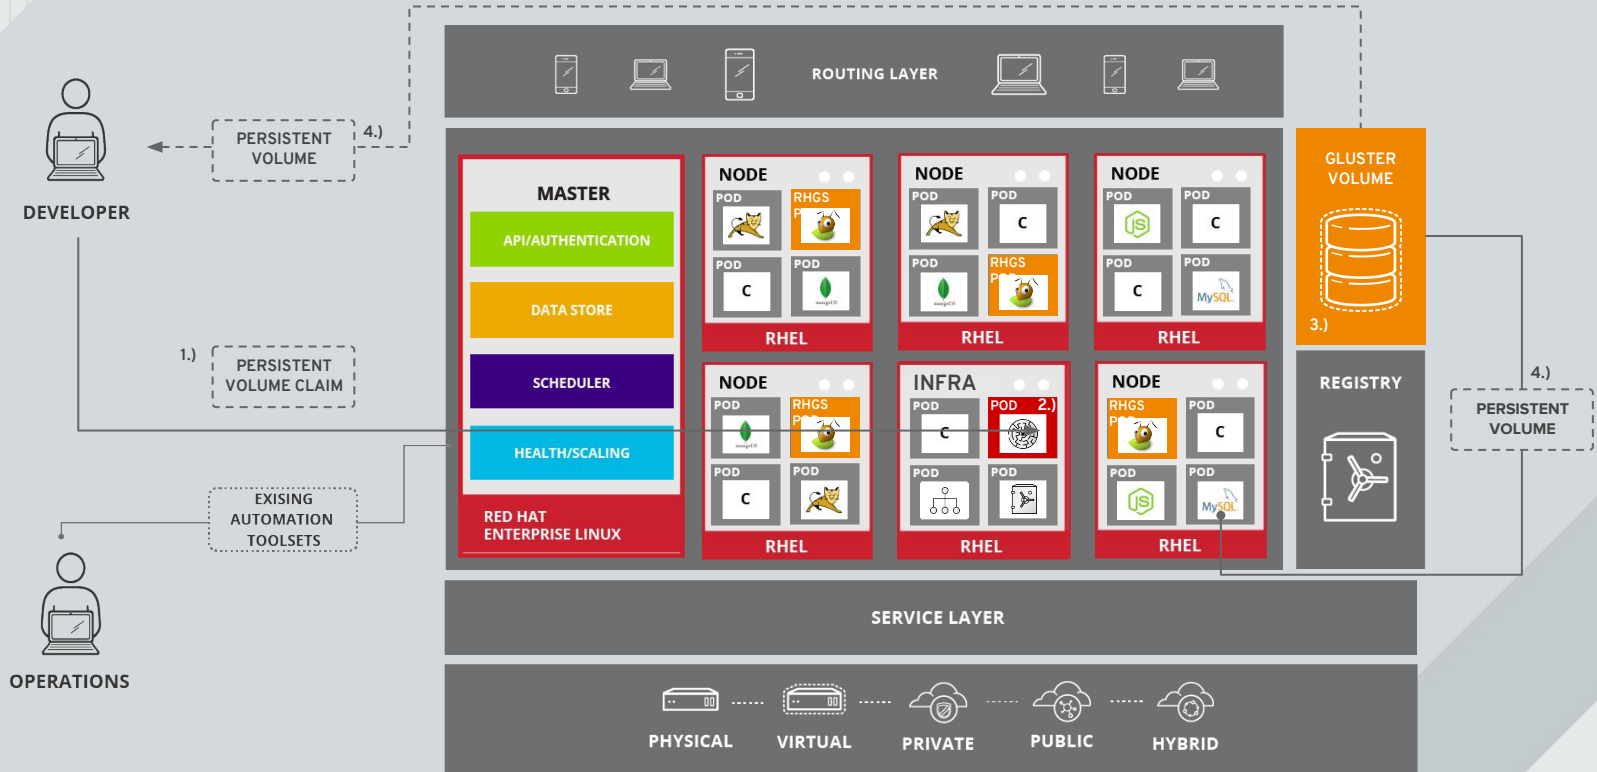
CONTAINER-NATIVE STORAGE ON OPENSHIFT



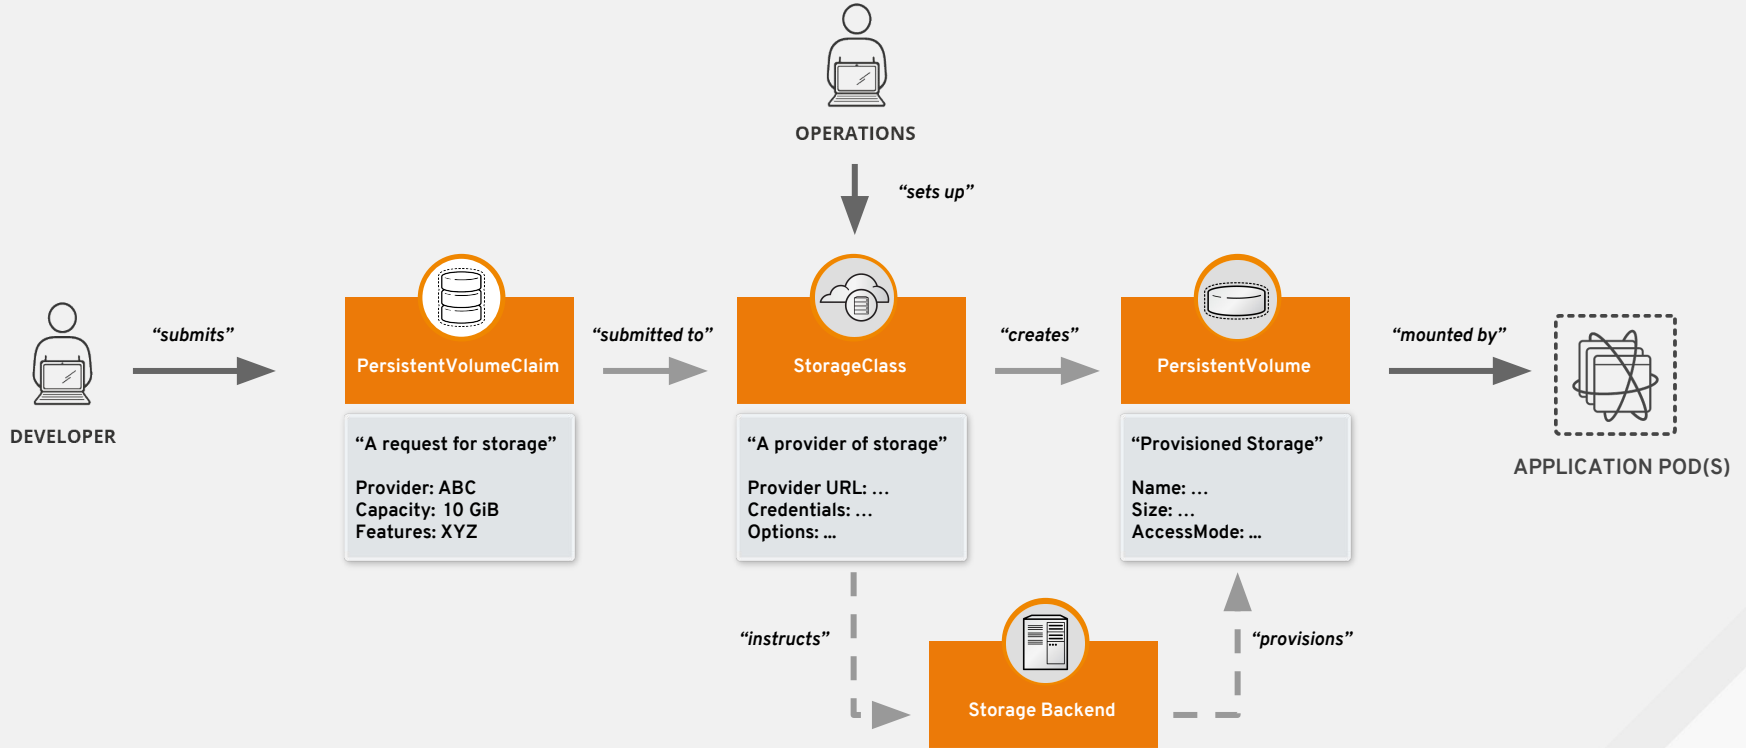
CONTAINER-NATIVE STORAGE ON OPENSHIFT



CONTAINER-NATIVE STORAGE ON OPENSHIFT

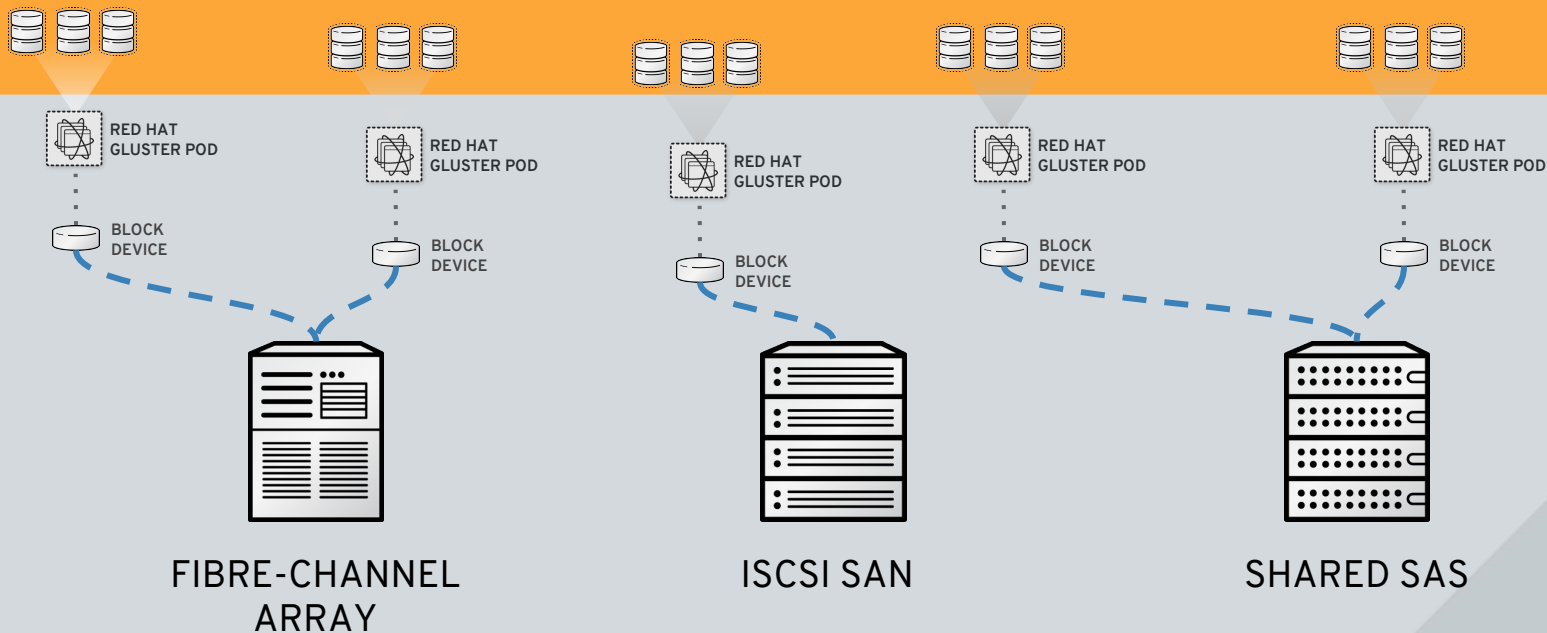


OpenShift PersistentStorage System



TRANSFORM LEGACY STORAGE TO CONTAINER READY

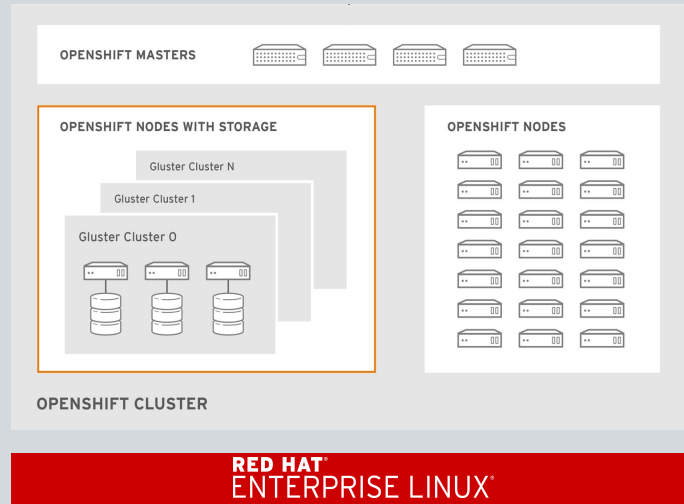
PERSISTENT VOLUMES PROVIDED BY CONTAINER-NATIVE STORAGE



SUMMARY

With Container-native Storage

- **Scalable**
(1000+ volumes)
- **Highly-Available**
(across availability zones)



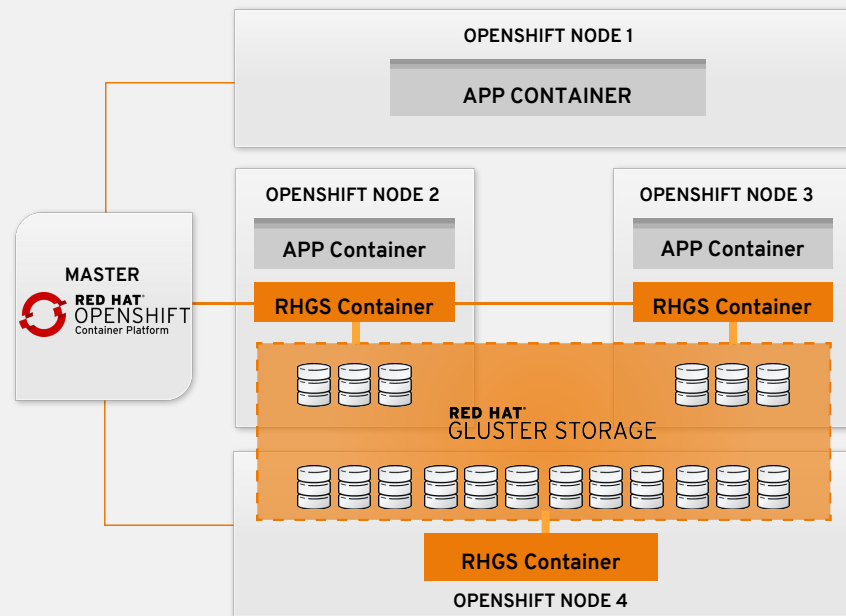
- **Automated**
(Dynamic Provisioning)
- **Integrated**
(installs with / runs on OpenShift)

Persistent, resilient and elastic storage...

... that travels with the platform.

Automated Container Native Storage deployment with OCP Advanced Installation

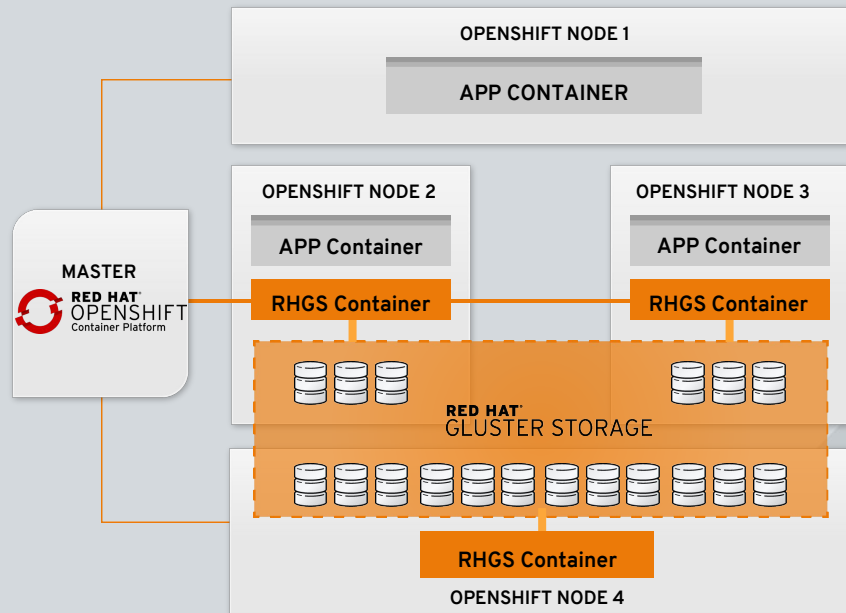
- OCP + CNS deployed as one cluster
- CNS with Block & File provisioners deployed
- OCP Registry deployed on CNS
- Ready to deploy Logging, Metrics on CNS
- CNS uninstaller



Custom Volume Naming

Previously PV Names (*vol_<UUID>* , *vol_1213456*)

- Specify new attribute in CNS storage class called '**volumenameprefix**'
- CNS backend volumes will be named ***myPrefix_NameSpace_PVClaimName_UUID***
- Easy to recognize, users follow naming convention,
- Easy to Search & Apply Policy based on prefix, Namespace, Project Name, or Claim Name



End-to-End CNS pv expansion

Can be done online from OCP

- Previously only available from Heketi CLI
- User edits PVC for the new size, triggering PV resize
- Fully Qualified for glusterFs backed PV's
- *Gluster-block PV resize will be added with RHEL 7.5*

How it Works/Example:

- Add to storage class AllowVolumeExpansion=true
- oc edit pvc claim-name
- Edit the field 'spec' → requests → storage: new value'

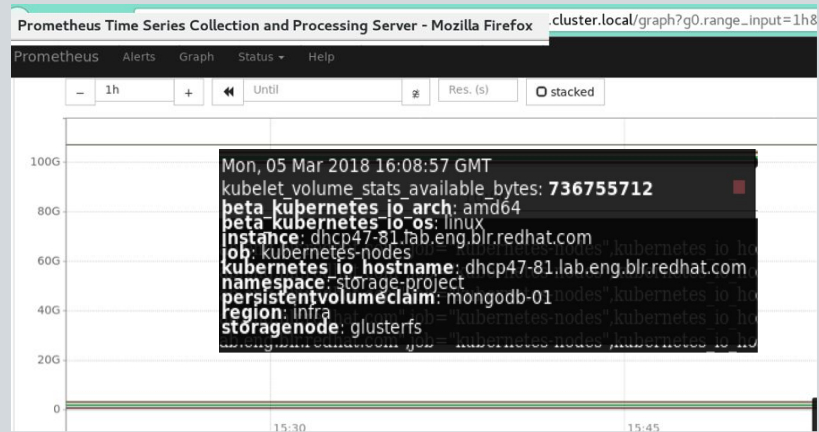
```
# Please edit the object below. Lines beginning with a pound sign and an empty line will abort the edit.
# If you open this file in a editor, it will be reopened with the relevant failures.
#
apiVersion: v1
kind: PersistentVolumeClaim
metadata:
  annotations:
    pv.kubernetes.io/bind-completed: "yes"
    pv.kubernetes.io/bound-by-controller: "yes"
    volume.beta.kubernetes.io/storage-class: glusterfs
    volume.beta.kubernetes.io/storage-protection: "true"
  creationTimestamp: 2018-02-07T12:03:36Z
  finalizers:
  - kubernetes.io/pvc-protection
  name: claim1
  namespace: default
  resourceVersion: "396"
  selfLink: /api/v1/namespaces/default/persistentvolumeclaims/claim1
  uid: ed8d425a-0bfe-11e8-b140-c85b7636c000
spec:
  accessModes:
  - ReadWriteMany
  resources:
    requests:
      storage: 12Gi
  volumeMode: Filesystem
  volumeName: pvc-ed8d425a-0bfe-11e8-b140-c85b7636c000
status: {}
```


Capacity usage metrics

CNS GlusterFS extended to provide PV volume metrics (including consumption) through Prometheus or Query

- Metrics available from PVC end point
- User can now know PV size allocated as well as consumed and use resize (Expand) of PV if needed from OCP
- Example Metrics added
 - `kubelet_volume_stats_capacity_bytes`
 - `kubelet_volume_stats_inodes`
 - `kubelet_volume_stats_inodes_free`
 - `kubelet_volume_stats_inodes_used`
 - `kubelet_volume_stats_used_bytesetc`

Prometheus



'curl'

```
# TYPE kubelet_volume_stats_available_bytes gauge  
kubelet_volume_stats_available_bytes{namespace="default",persistentvolumeclaim="claim1"} 8.543010816e+09  
# TYPE kubelet_volume_stats_capacity_bytes gauge  
kubelet_volume_stats_capacity_bytes{namespace="default",persistentvolumeclaim="claim1"} 8.57735168e+09
```

Heketidb - maintenance

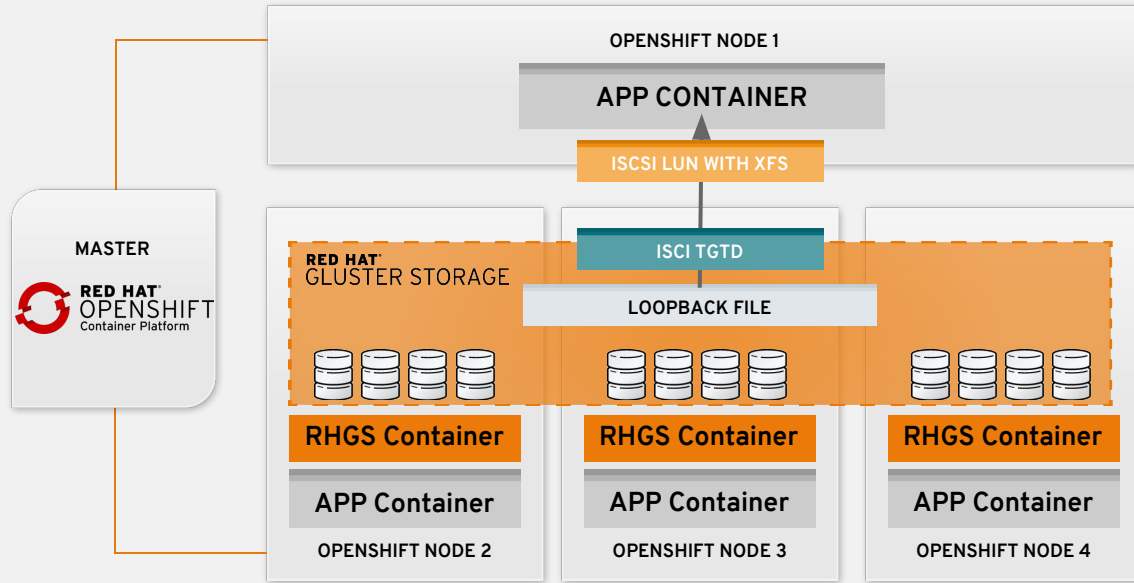
```
$ heketi-cli db dump
```

```
$ heket-cli db import  
(can even be done without heketi daemon running)
```

```
$ heketi-cli db repair
```


CONTAINER-NATIVE STORAGE BLOCK STORAGE

ENABLES LOGGING/METRICS ON CNS




BRICK-MULTIPLEXING

BEFORE



300 MB RAM
/Brick

AFTER

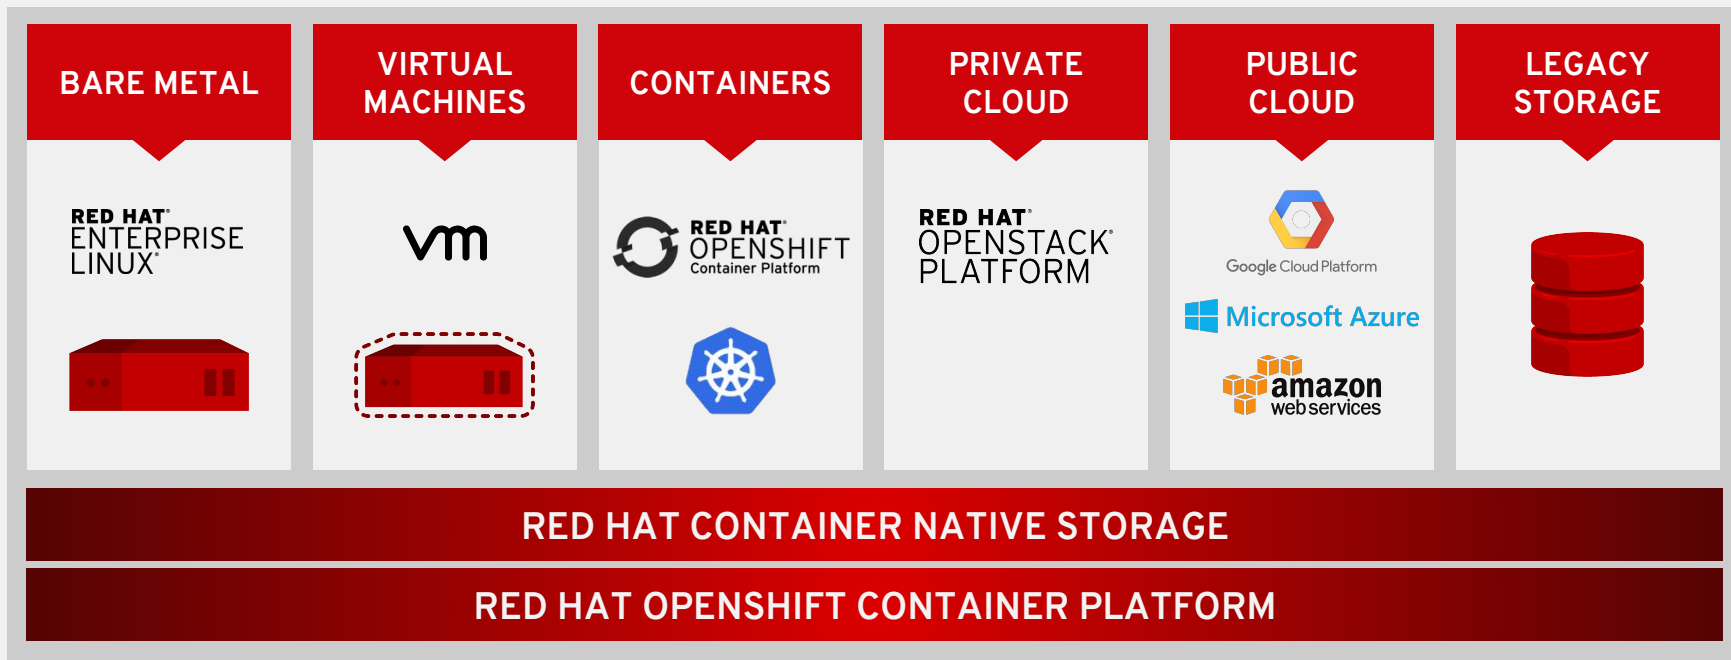


30 MB RAM
/Brick



Recommended 8GB RAM baseline + 30MB per brick

Consistent Storage Experience Across The Hybrid Cloud

Application Portability And Lower Costs



Simplify Container Availability



AVAILABILITY
ZONE A



AVAILABILITY
ZONE B



AVAILABILITY
ZONE C

Red Hat Ceph Storage

- **RED HAT®**
CEPH STORAGE

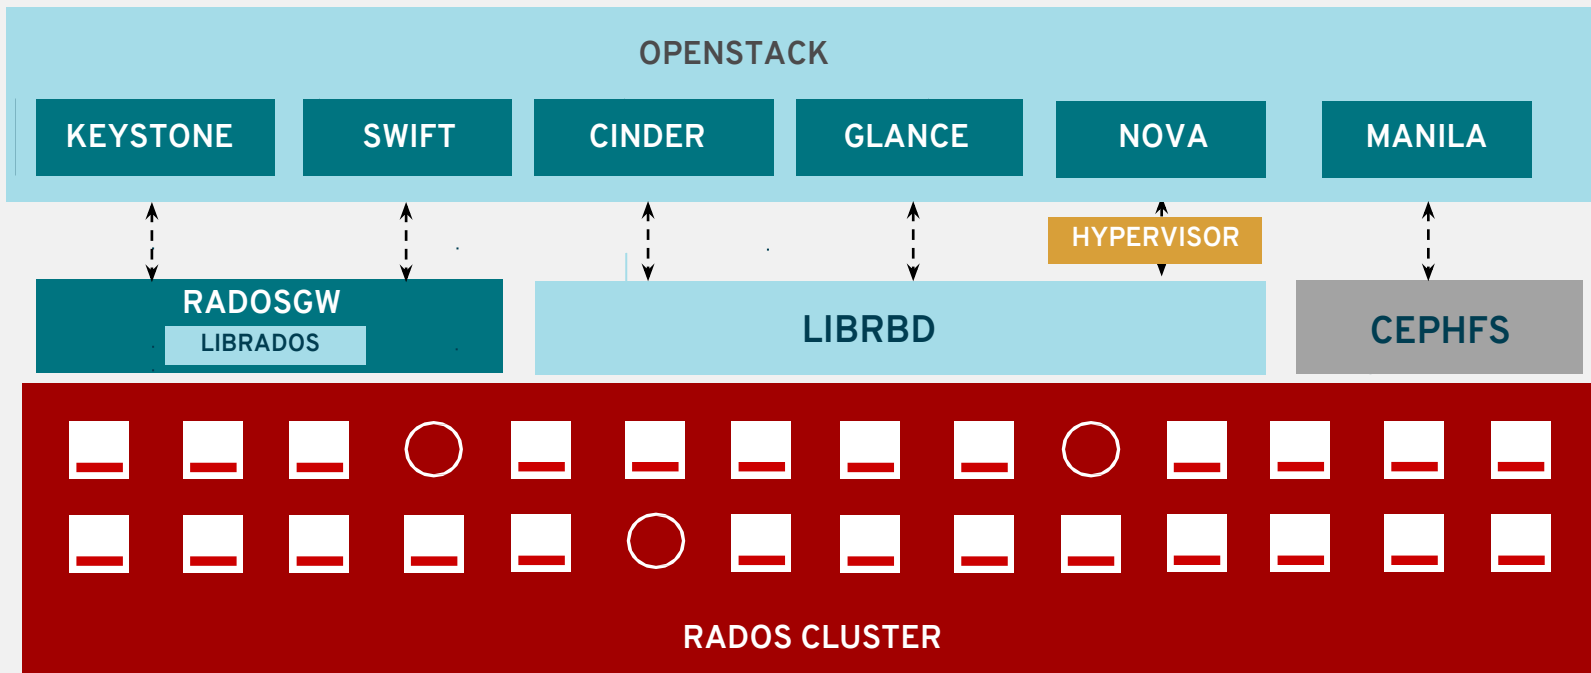
OpenStack Storage (Cinder/Glance/Swift/Nova/Manila)

S3 scale-out Object Storage (AWS Compatible S3)


Elastic Data Lake Storage (S3A)

Easy tools to ingest data, NFS gateway for RADOS-GW

RED HAT CEPH STORAGE AND OPENSTACK



Which OpenStack Cinder driver are you using?



Red Hat Ceph Storage

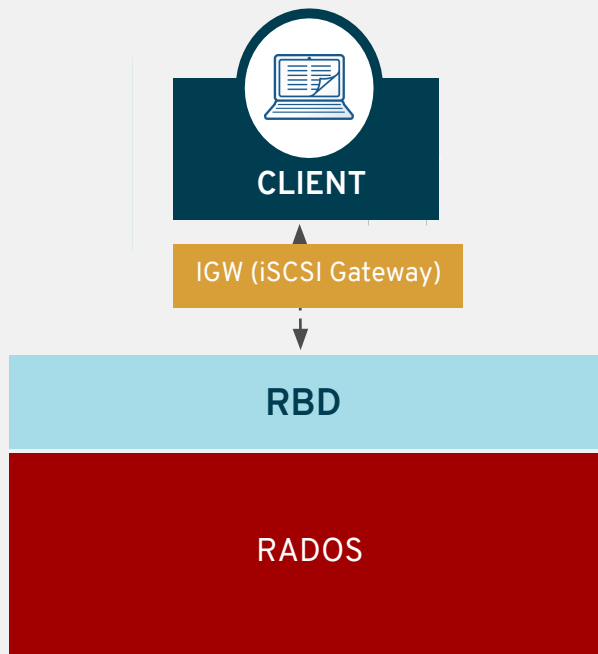
What's new...

NFS RGW Gateway



- Provides NFS file access on top of RGW object API
- NFS V4 support
- Hosts can mount and access the object namespace using NFS mount
- Files are internally stored as objects inside buckets
- Gateway seamlessly translates between host file semantics and RHCS object semantics.
- No caching of files or data required

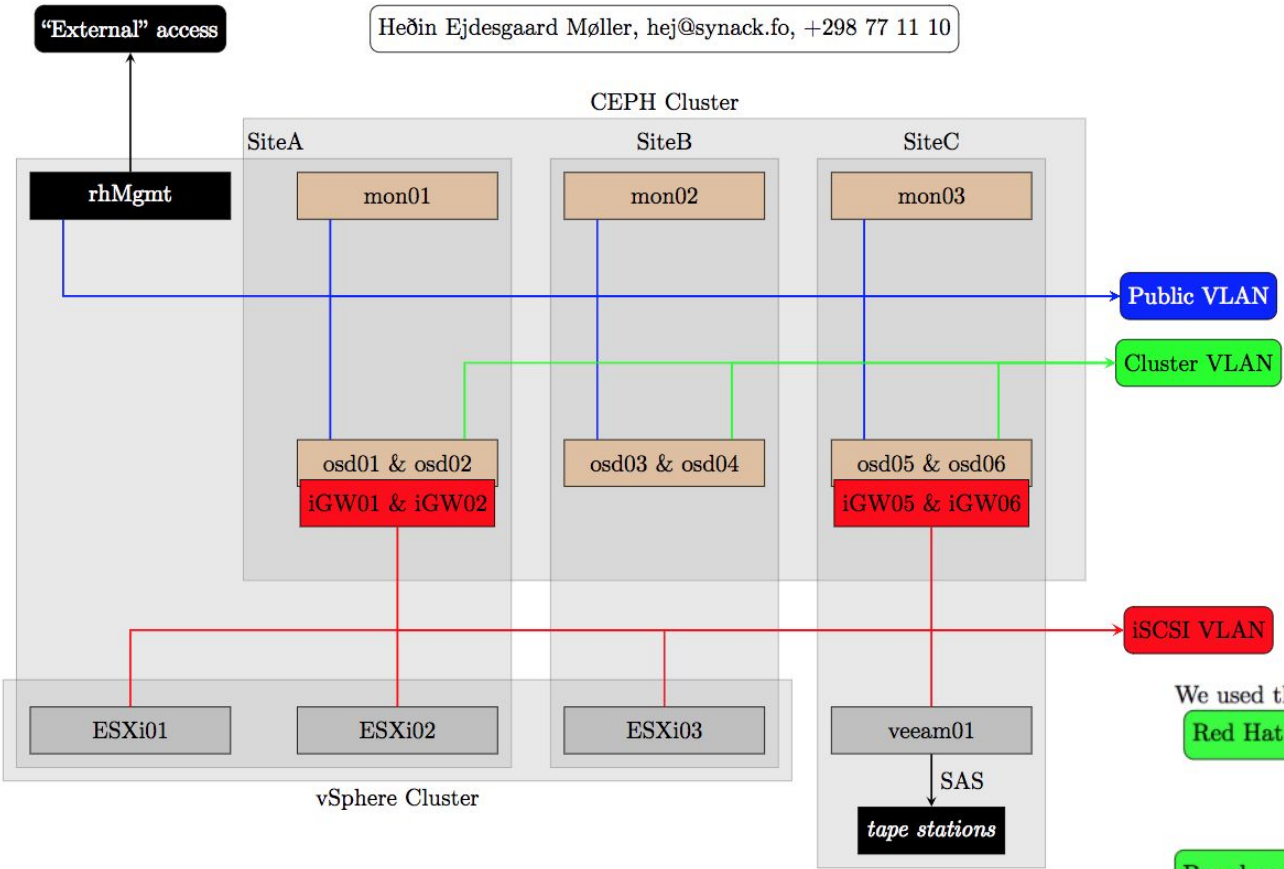
ISCSI BLOCK STORAGE



Sending a lifeline to legacy

- Block-level access via industry storage networking standard to storage devices and data on VMware, RHV, native Windows & RHEL
- Backwards compatibility facilitating transition between architectures
- Ceph for backup and recovery ancillary use cases
- ROI benefits from consolidated, unified storage platform
- Complement to NFS for migrating legacy workloads

Heðin Ejdesgaard Møller, hej@synack.fo, +298 77 11 10



We used the following components.

Red Hat Ceph Storage 3

VMware ESXi-6.5 & VCSA-6.5

Based on Luminous-12.2.x

Red Hat Ceph Storage - now containerized

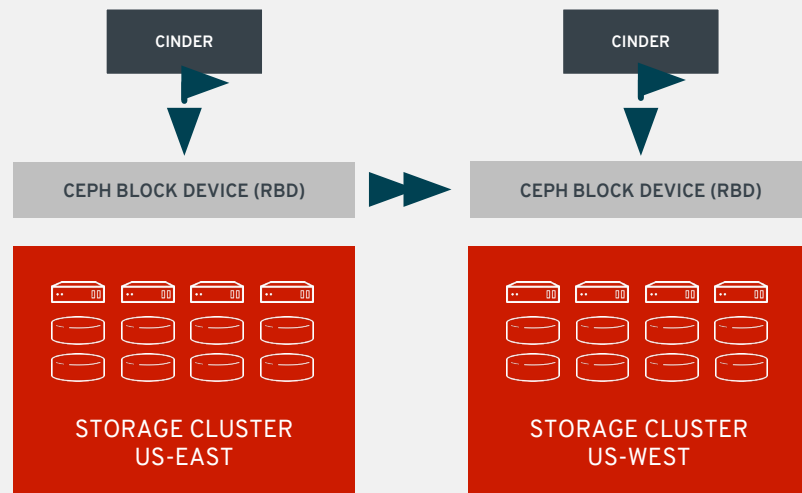
Deploy Red Hat Ceph enterprise storage in Linux containers

for simplified operations and a smaller hardware footprint. Containerized storage daemons enable users to run Red Hat Ceph Storage on fewer servers by co-locating services that previously required dedicated hardware while avoiding the risk of resource conflicts. Preliminary tests based on a standard Red Hat Ceph Storage cluster configuration showed lowered hardware expenditure by at least 24 percent.

RBD MIRRORING

New daemon 'rbd-mirror' synchronises Ceph images from one cluster to another

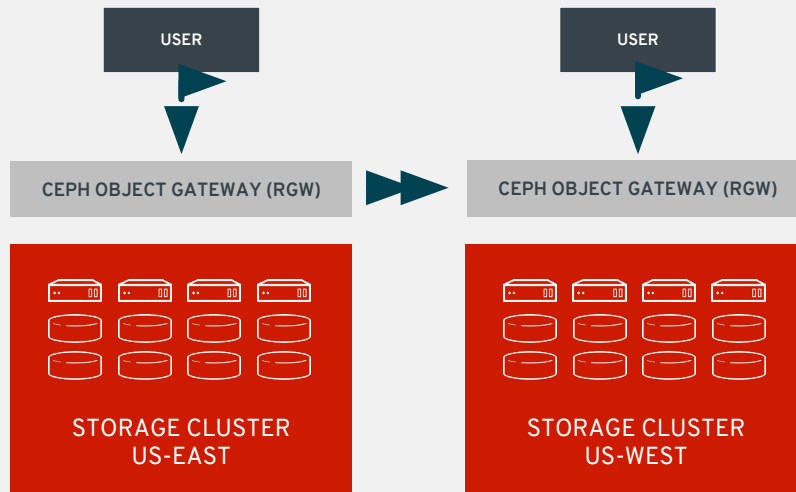
- Relies on two new RBD image features:
 - Journaling: enables journaling for every transaction on the image
 - Mirroring: tells the rbd-mirror daemon to replicate images
- Images have states: primary and non-primary (promote and demote calls)



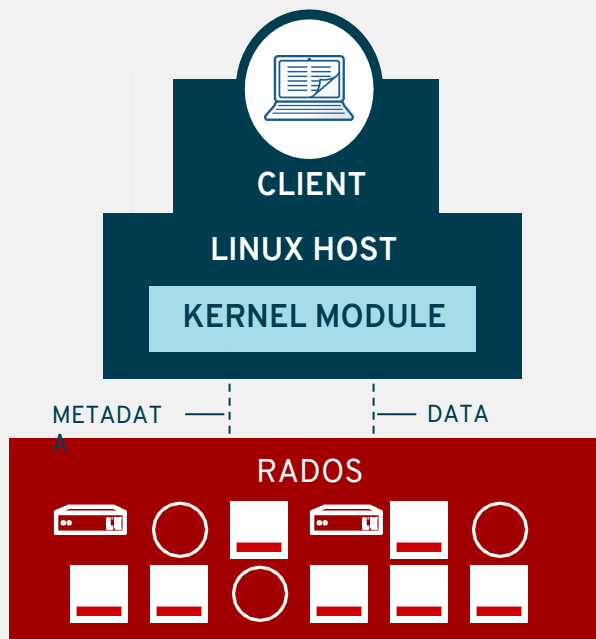
Rados Gateway multisite

Global object storage clusters with a single namespace

- Enables deployment of clusters across multiple geographic locations
- Clusters synchronize, allowing users to read from or write to the closest one



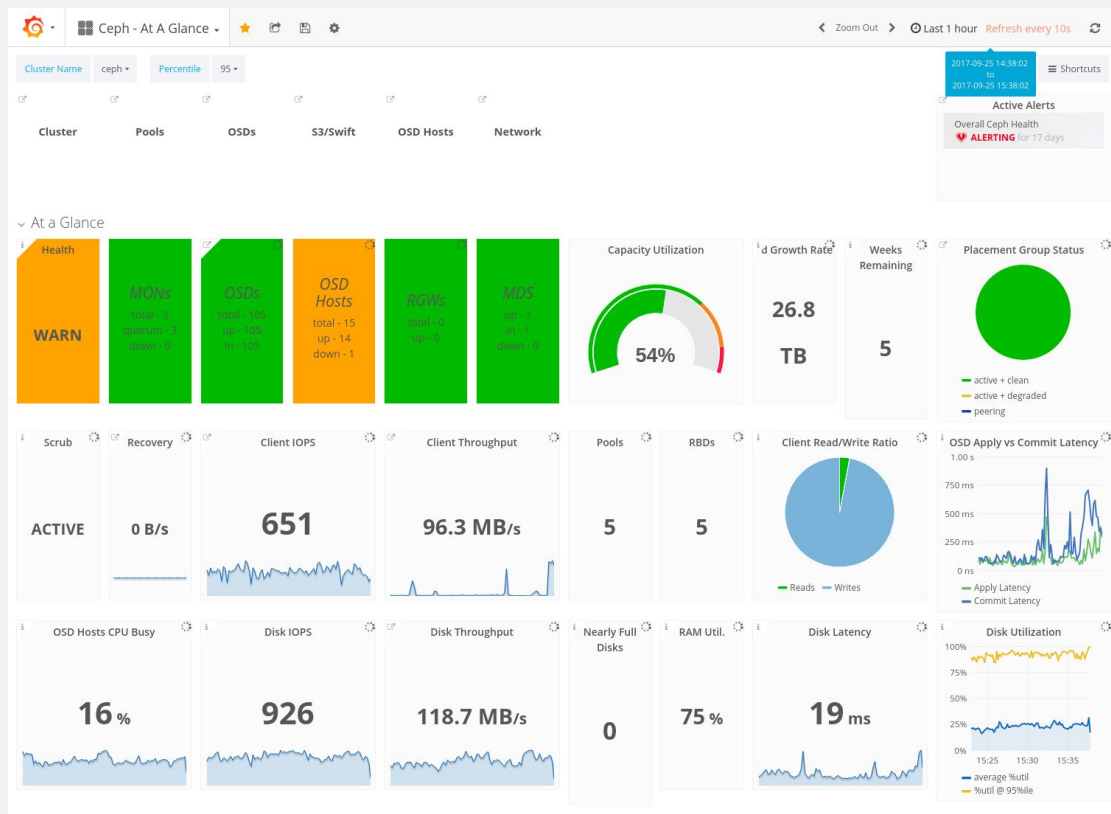
FILE STORAGE with CEPHFS



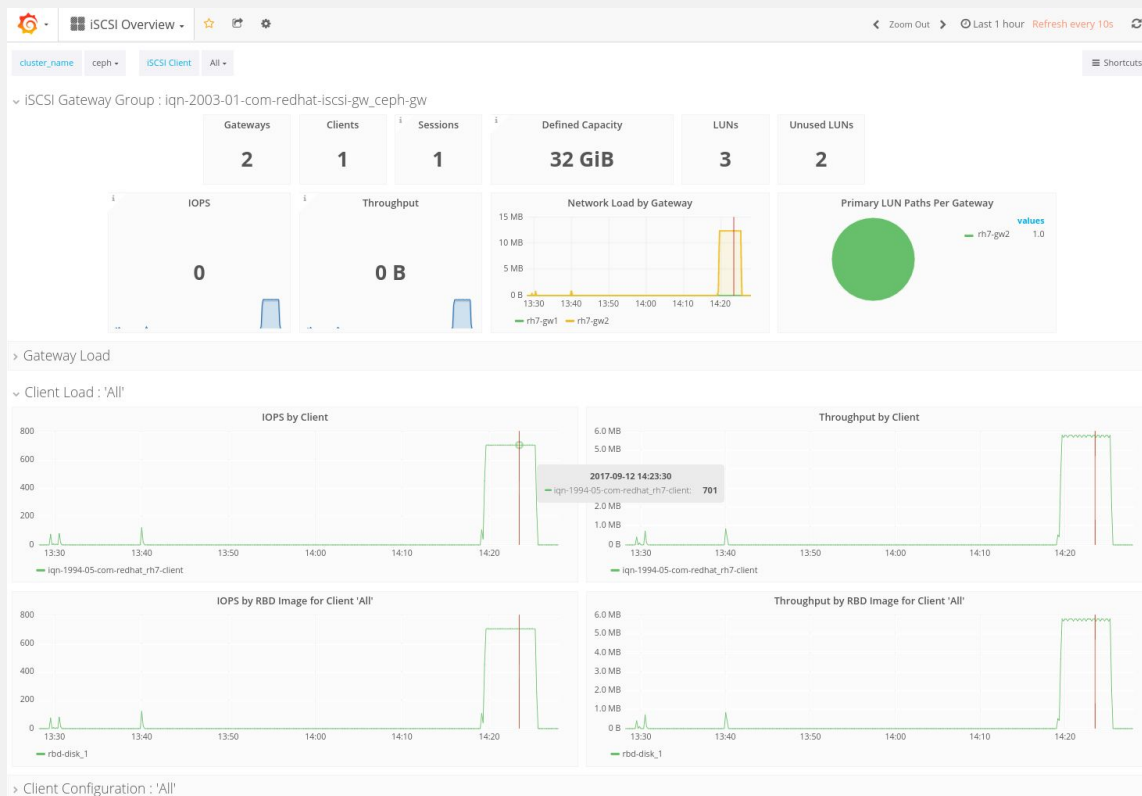
File services at scale

- Sharing infrastructure with RBD, RGW deployments.
- Very high scale data throughput (RADOS). One of only FS in the world where you can put PB's of data.
- Scale-out metadata performance (multiple metadata data servers).
- OpenStack Manila integration, nfs-ganesha integration (tech-preview)

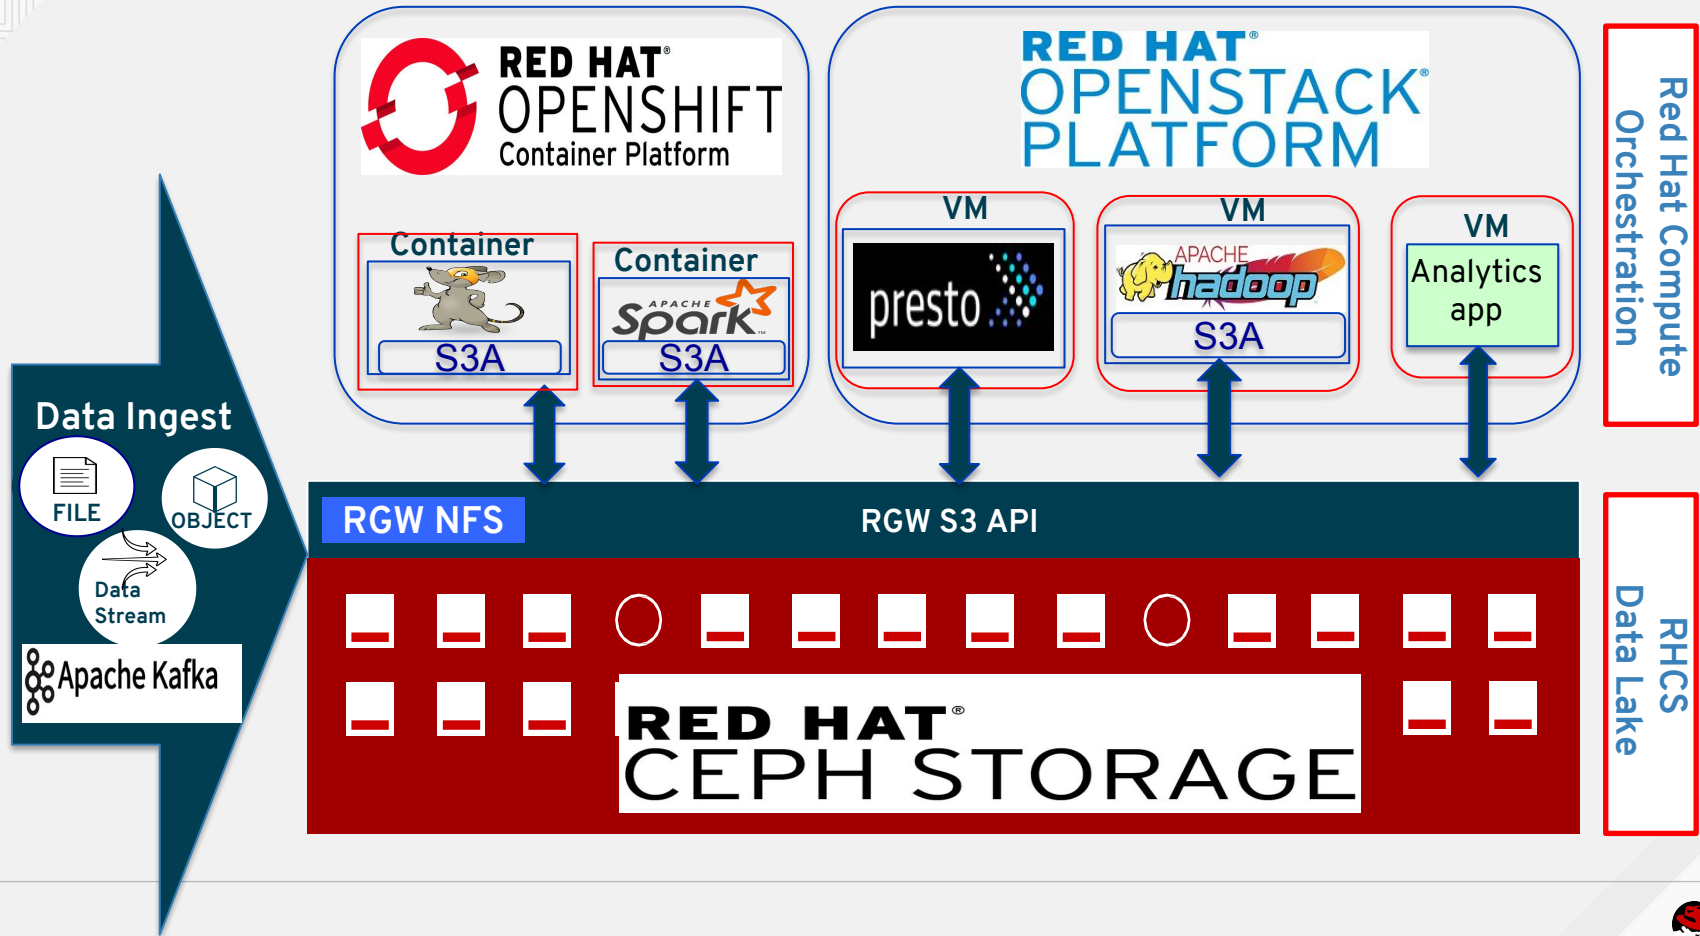
THE AT-A-GLANCE DASHBOARD




UNDERSTANDING LOAD - “cephmetrics”

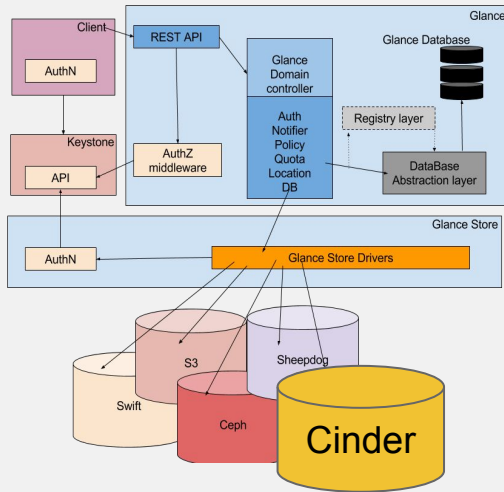


RHCS Data lake vision – On-demand compute and centralized object storage





Red Hat Ceph Storage integration with OpenStack




Create a volume from an image with a cinder backend.

Capacity Derived IOPs


Values can be scaled (ie multiplied) by volume size for setting dynamic limits.

QoS Spec Key	QoS Spec Value	2 GB Volume	5 GB Volume
Read IOPS / GB	10000	20000 IOPS	50000 IOPS
Write IOPS / GB	5000	10000 IOPS	25000 IOPS

Red Hat Hyperconverged Infrastructure RHHI



Traditional Architecture

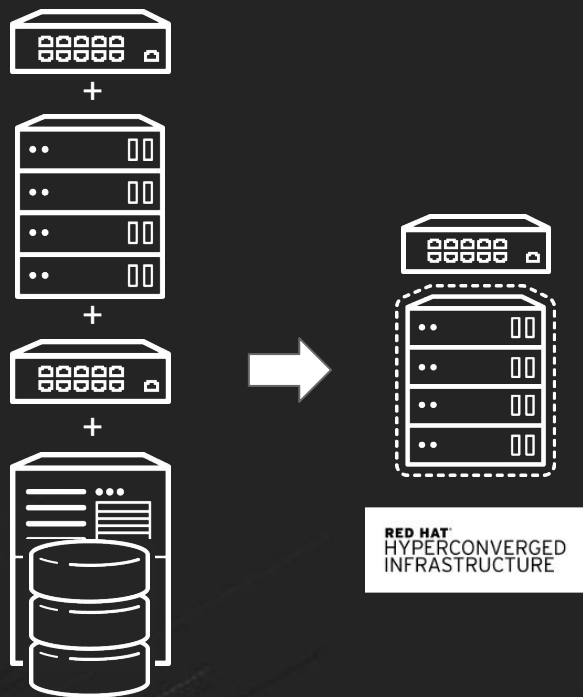


Converged Infrastructure



Hyper converged Infrastructure

INFRASTRUCTURE CONSOLIDATION & OPERATIONAL EFFICIENCY






Traditional Architecture

- Single budget for compute & storage
- One team to managing infrastructure
- Simplified planning & procurement
- Streamlined deployment & management
- Single support stack for compute and storage
- Maximizing resource utilization

RED HAT HYPERCONVERGED INFRASTRUCTURE

OPTIMIZE, INTEGRATE, MANAGE

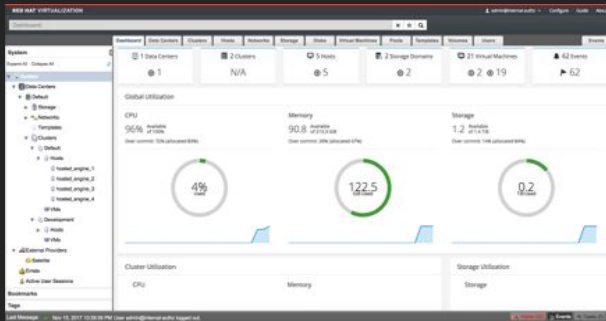
VM VM VM VM VM VM VM VM VM VM



RED HAT
HYPERCONVERGED
INFRASTRUCTURE

MANAGING YOUR RED HAT HYPERCONVERGED INFRASTRUCTURE INSTALLATION

RED HAT VIRTUALIZATION MANAGER




Provides secure web interface (HTTPS) as well as role based access for administration of compute, network, and storage resources.

RED HAT CLOUDFORMS (optional)



Provides orchestration, self-service catalogs, central management for many Red Hat Hyperconverged Infrastructure deployments, chargeback, and policy enforcement.*

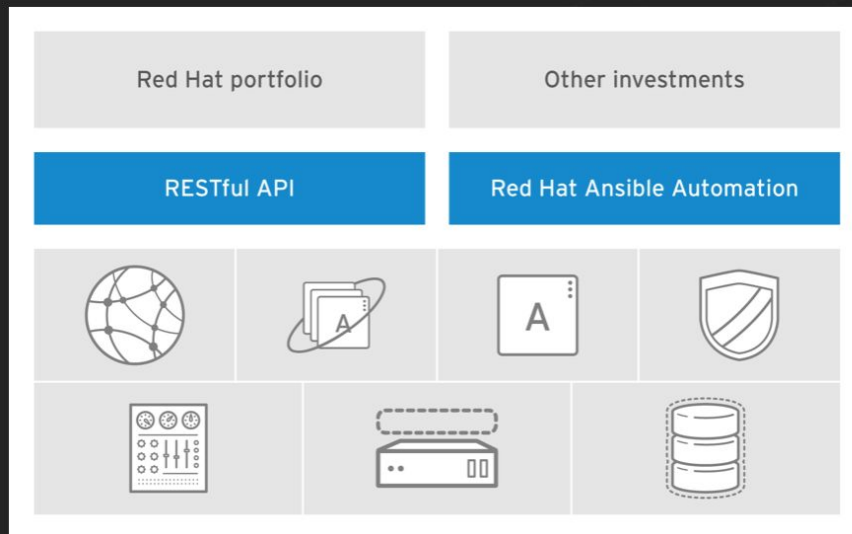


*CloudForms is not part of the Red Hat Hyperconverged Infrastructure product or SKU; it is a separate product.

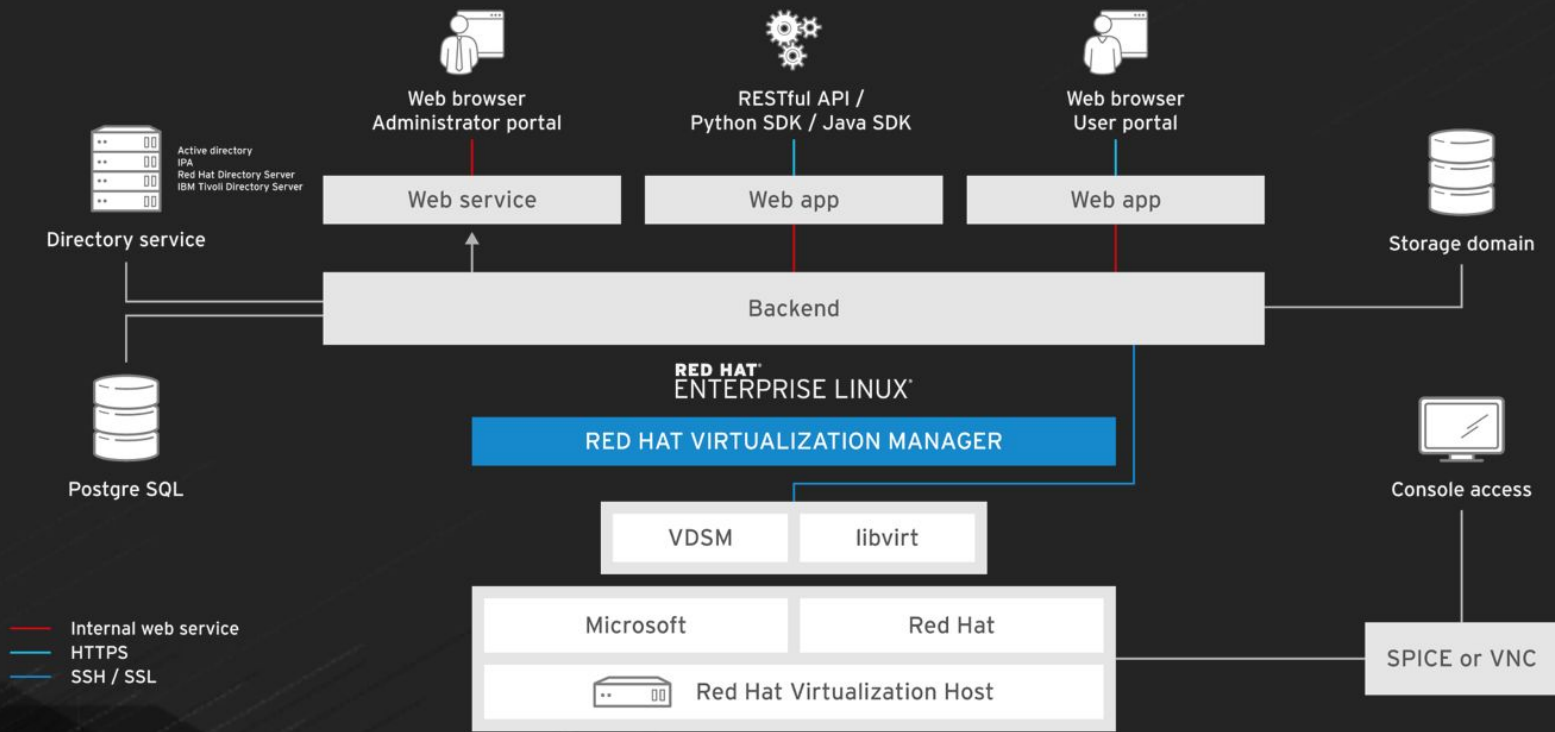
RED HAT VIRTUALIZATION

RED HAT VIRTUALIZATION OVERVIEW

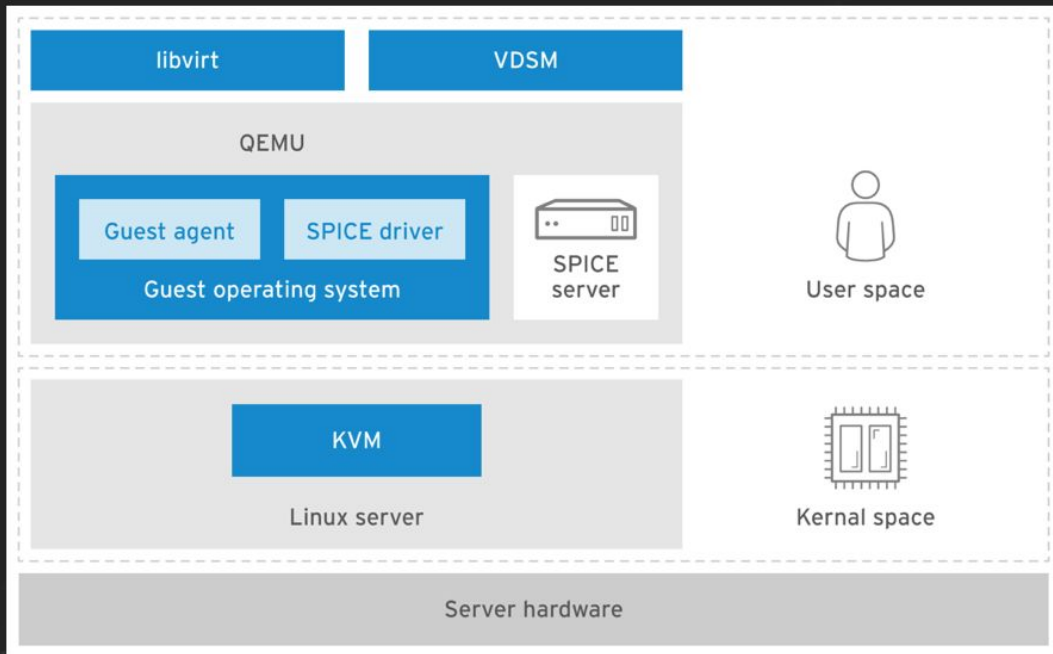
- ▶ **CENTRALIZED MANAGEMENT** for the KVM hypervisor as well as compute, network, and storage resources
- ▶ **ENTERPRISE FEATURES** to support business-critical applications
- ▶ **CROSS-PORTFOLIO** integration, APIs, and SDKs to enable automation
- ▶ **RED HAT VIRTUALIZATION** is built on Red Hat Enterprise Linux+KVM



OVERVIEW OF RED HAT VIRTUALIZATION



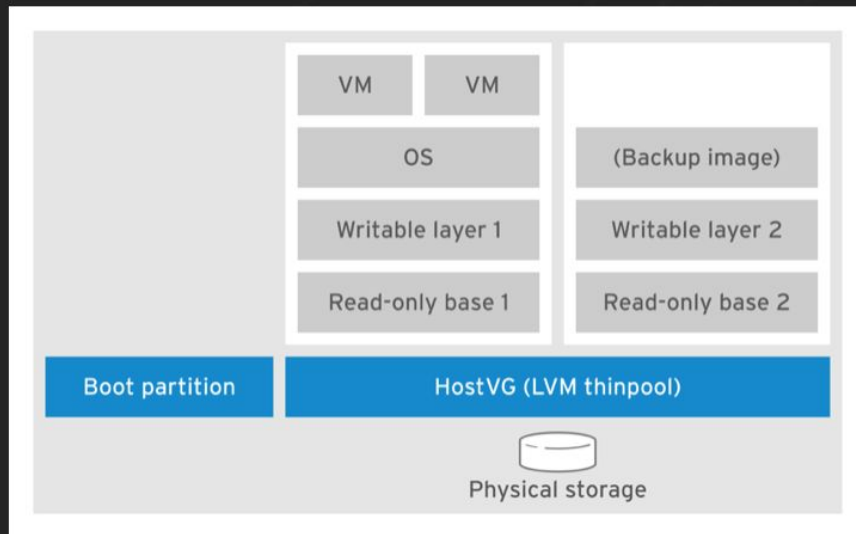
KERNEL-BASED VIRTUAL MACHINE (KVM)



RED HAT VIRTUALIZATION HOST

LIGHTWEIGHT HOST

- ▶ **RED HAT VIRTUALIZATION HOST**—Purpose built node built on Red Hat Enterprise Linux
- ▶ **CAN BE DEPLOYED** via ISO, PXE, USB, cloned
- ▶ **WRITABLE** root file system
- ▶ **USES TRIMMED-DOWN** Anaconda installer
- ▶ **COCKPIT** administrative console
- ▶ **SECURITY AND SERVICES** are pretuned to support virtual machines

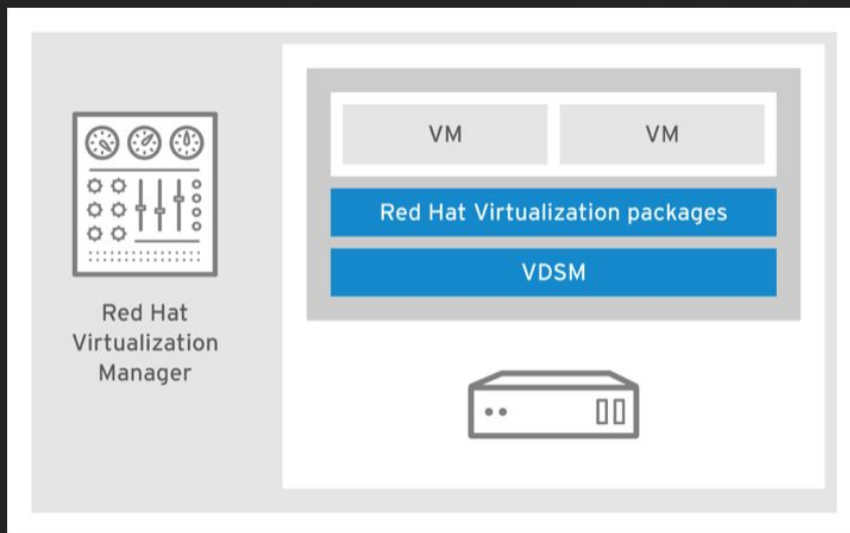


Red Hat Virtualization Host is designed around LVM Thinpools and “imgbased,” resulting in a lightweight and flexible architecture.

RED HAT ENTERPRISE LINUX NODE

FULL HOST

- ▶ **RED HAT VIRTUALIZATION 4** supports Red Hat Enterprise Linux 7 as a node
- ▶ **USES QEMU-KVM-Red Hat Virtualization**
- ▶ **LARGER FOOTPRINT** as compared to Red Hat Virtualization Host
- ▶ **RED HAT VIRTUALIZATION MANAGER** will configure security and VDSM
- ▶ **COCKPIT** needs to be manually installed and configured



Red Hat Enterprise Linux 7 is fully supported as a host in Red Hat Virtualization. Red Hat Virtualization-related packages and policies are deployed by Red Hat Virtualization Manager.

RED HAT HYPERCONVERGED INFRASTRUCTURE 1.0

Basic features



Single point of management for
virtual resources

CPU pinning

RBAC and tiered access

Power management

VM templates

Firewall/SELinux

Support for Red Hat

Enterprise Linux and Windows

HA virtual machines

Secure Browser based management

RESTful API

Python, Ruby, and Java SDKs

Live migration

RED HAT HYPERCONVERGED INFRASTRUCTURE 1.0

Advanced features



Host affinity and anti affinity

Migrate and import VMs

Automated resource mgmt

Load balancing

CPU QoS

Hot-add memory and CPU

Resource reservation

Overcommit (memory ballooning)

Memory page sharing

Large-page support

Import VMs from VMware

RED HAT HYPERCONVERGED INFRASTRUCTURE 1.0

Network features



VLAN tagging

Network QoS

NIC bonding

VM-FEX support

Open virtual network (tech preview)

IPv6 support (guest)

Jumbo frames

Network labels

RED HAT HYPERCONVERGED INFRASTRUCTURE 1.0

Storage features



Geo-replication

Sharding support

Live snapshots and merge

Block discard

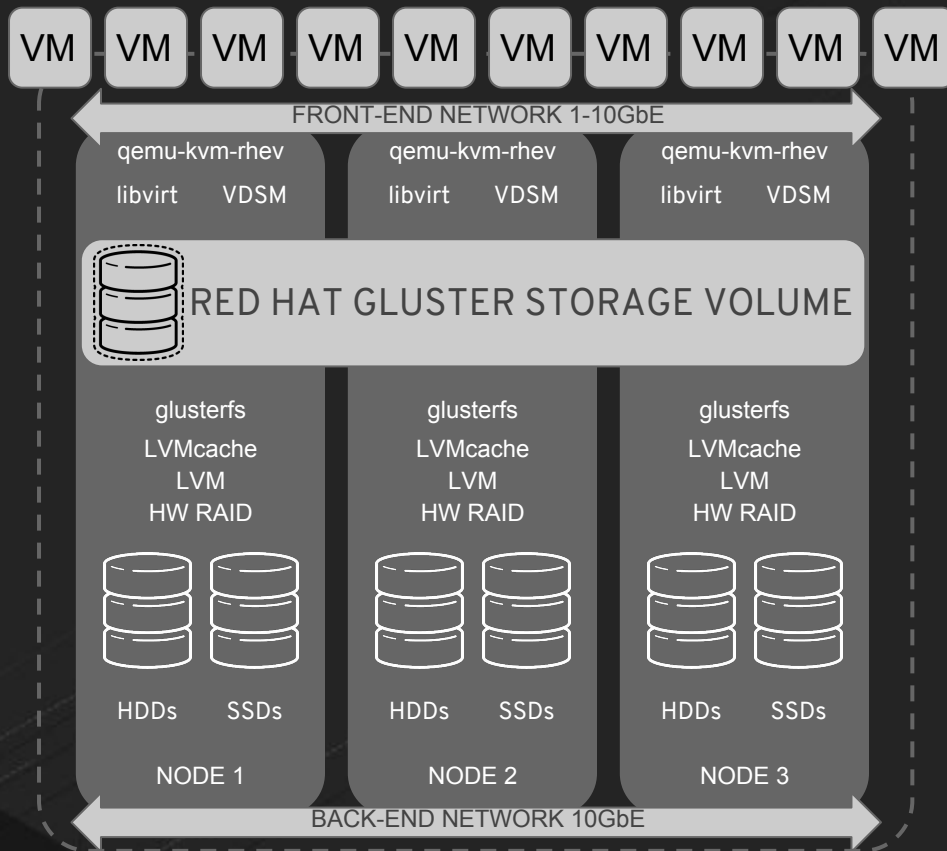
RESTful API for backup/restore

3-way Data Replication

Thin and thick provisioning

Storage-based fencing

RED HAT HYPERCONVERGED INFRASTRUCTURE ARCHITECTURAL VIEW - POD



3,6 or 9 Pod Configuration

RED HAT HYPERCONVERGED INFRASTRUCTURE OR RHV + RHGS?



RESOURCES

Red Hat Hyperconverged Infrastructure Documentation

<https://access.redhat.com/products/red-hat-hyperconverged-infrastructure>

All Red Hat Virtualization Documentation

<https://access.redhat.com/documentation/en/red-hat-virtualization>

All Red Hat Gluster Storage Documentation

<https://access.redhat.com/documentation/en/red-hat-storage>

Questions ?



THANK YOU



plus.google.com/+RedHat



facebook.com/redhatinc



linkedin.com/company/red-hat



twitter.com/RedHatNews



youtube.com/user/RedHatVideos